ATTENTION & THE INDETERMINACY OF VISUAL EXPERIENCE

by

JAMES DAVID STAZICKER

A dissertation submitted in partial satisfaction of the requirements for

the degree of

Doctor of Philosophy

in

Philosophy

in the Graduate Division of the

University of California, Berkeley

Fall 2011

Committee in charge

Professor John Campbell (Co-Chair)

Professor Alva Noë (Co-Chair)

Professor Lucia Jacobs

Abstract

Attention and the Indeterminacy of Visual Experience

by

James David Stazicker

A dissertation submitted in partial satisfaction of the requirements for the degree of Doctor of Philosophy in Philosophy, in the Graduate Division of the University of California, Berkeley

Committee Co-Chairs: Professor John Campbell and Professor Alva Noë

Visual attention plays an important role in epistemology. You come to know about things you see by paying attention to them, and thereby taking them up in thought and belief. Recent work in the cognitive sciences shows that attention also has a quite different role. Attention dictates not only what you take up in thought and belief, but also how things appear to you in visual experience. Attention enhances the neural signal that's processed for attended aspects of a scene, and so highlights those aspects by changing their visual appearance. For example, drawing attention to a colour makes it appear more saturated; drawing attention to a shape makes it appear larger. This raises a challenge for the epistemology of attention. It threatens to show that attention is a systematic source of illusion, rather than a reliable source of knowledge. Indeed, some scientists working in the area draw that conclusion.

To meet this challenge, I argue that visual experience is experience of *determinable* properties: properties which admit of more specific determinations, as *red* is determined by *crimson*, and *90 to 110 feet long* is determined by *100 feet long*. One determinate property determines many determinable properties. So one determinate shape or colour may take on different appearances, in veridical experiences of its different determinables. And one determinable property has many determinates. So different determinate properties may share an appearance, in veridical experiences of one determinable which they all determine.

I argue that the experimental data about attention and visual appearance are in fact well analysed in this way. For example, attention changes the appearance of a colour, giving it the same appearance as a more saturated but unattended colour; visual appearances remain veridical through this change, because they consist in experiences of

different determinable colours, each of which the stimulus really instantiates. More generally, I argue that we need to recognise the role of determinable properties in vision, if we're to understand what experiments tell us about the contents of visual experience.

To show that visual experience is experience of determinable properties, I draw on empirical work about the limits of visual resolution. Visual processes are not sensitive to the finest details of a scene; they're sensitive to determinable properties, not maximally determinate properties. I argue that visual experience is experience of a property only where visual processes are sensitive to that property. I explore the metaphysics of determinable properties, arguing that they have a natural unity which merely disjunctive properties lack. On this basis, I show how visual experience of a determinable property has the phenomenological unity characteristic of experience of a single shape or colour.

With this account of visual indeterminacy in hand, we can also re-assess traditional assumptions about the nature of attention, and about the relationship between attention and visual experience. In traditional Anglophone philosophy, attention is often conceived as a window onto visual experience: attention gives you access to the contents of visual experience, but does not alter them. While this is clearly untenable in light of the empirical findings, a weaker view has attracted recent interest: visual experience is *constitutively* independent of attention; your seeing as you do does not consist partly in your attending. I make the relevant notion of constitutive dependence precise, and criticise Ned Block's argument for a version of this view, on the grounds that he underestimates the indeterminacy of visual experience.

I then argue that, in some instances, visual experience does depend constitutively on attention – and also *vice versa*. I defend William James's definition of conscious attention as the 'focalization, concentration of consciousness'. So conceived, I propose, attention is not a further mode of consciousness, over and above perceptual experience and thought. Rather, attention consists in a focusing of these modes of consciousness. For example, experiments have shown that attention increases visual resolution: when you attend, you see more determinate properties. Here, I argue, attention and visual experience are mutually constitutive: they consist in one another. More generally, I show how to understand attention in terms of the focusing of conscious cognition and perception. This, I propose, is the form of conscious attention we exploit when we take up what we see in thought and belief, and thereby come to know about it. I conclude by exploring some consequences for intentionalism about visual experience.

Contents

I am especially grateful to the chairs of my dissertation committee, John Campbell and Alva Noë, and also to Mike Martin, for their thoughtful criticisms, warm encouragement, and hard work on my behalf. Heartfelt thanks to all three.

I learned a great deal from presenting work on this material on various occasions: a Fall 2008 meeting of Michael Silver's vision lab at Berkeley; Barry Stroud's Spring 2009 seminar and John Campbell's Fall 2009 seminar at Berkeley; a Fall 2009 meeting of Berkeley's Wollheim Society; the 2010 Berkeley-Stanford-Davis Graduate Student Philosophy Conference; the 2010 London-Berkeley Graduate Student Philosophy Conference; the EXRE / University of Fribourg conference on 'Phenomenal Presence' in June 2010; a Philosophy Department Colloquium at Berkeley in November 2011. Thanks to participants for discussion, and to Bert Baumgartner and Craig French for helpful comments.

Berkeley has been a wonderful place to do philosophy of psychology over the last few years. Too many people there – philosophers and scientists, faculty and students, the permanent and the passing-through – have contributed in one way or another to this dissertation for me to mention them all. But some conversations stand out vividly as having been especially helpful: with Joseph Barnes, Tony Bezsylko, Ned Block, Michael Caie, Tim Crane, Zoe Drayson, Bert Dreyfus, Peter Epstein, Hannah Ginsborg, Anjana Jacob, Mark Eli Kalderon, Ayelet Landau, Tamar Lando, Geoffrey Lee, Carlos Montemayor, Matt Parrott, Richard Price, Ariel Rokem, Sherry Roush, John Schwenkler, Michael Silver, Klaus Strelau, Barry Stroud and Seth Yalcin. Many thanks to them.

Chapter 1

INTRODUCTION

Studies of attention are a mainstay of contemporary cognitive science. Understanding the mechanisms of attention has proved to be essential for understanding a range of information-processes, not least the processes of vision. By comparison, and notwithstanding some recent work, attention has been a peripheral concern for philosophers. This is odd, because we need to understand attention if we're to understand two things with which philosophers have been very much concerned: the role of visual experience as a distinctive source of knowledge, and the nature of visual experience itself. In order to understand these things, I will argue, we need to understand how attention interacts with visual experience, and we need to understand what attention is.

Visual experience is an aspect of consciousness. It is the distinctively visual aspect of an overall conscious episode in virtue of which you see. There is an intuitive understanding of this aspect of consciousness according to which it is a way of being simply open to your environment – a way of simply representing, or being presented with, what is actually there around you. This intuitive idea will need refining later on, but for now we can work with it as stated.

For example, consider Figure 1 overleaf. It's natural to suppose that, in order to describe your experience of the figure, we should simply describe the figure itself. We should describe the overall pattern of the figure, and we should describe its more specific features: the orientations of the lines; the shapes of the triangle and the square; the locations and the colours of all these elements. Presumably, we should also include information about the spatial relations in which the figure stands to you, how well lit it is, and so on. Your experience of the figure will differ, depending on how far away and at what angle you hold it, and depending on the lighting under which you see it. Nonetheless, the idea is that by describing all these features of the figure, we can also describe your visual experience of it.

Figure 1

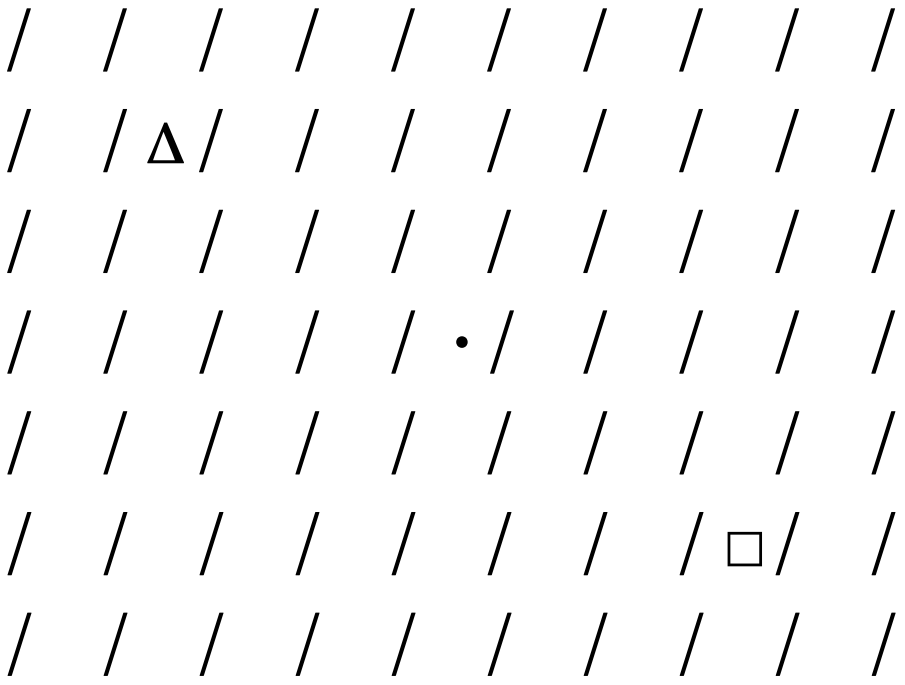

With respect to both the nature and the epistemic role of visual experience, there is a substantial theoretical advantage in this intuitive way of thinking. By understanding visual experience in this way, we can see how this aspect of consciousness might consist in, or at least be explained by, the gathering of information from your environment which is carried out by processes in the visual system. Equally, we can see how visual experience might be the distinctive source of knowledge which it seems *prima facie* to be. Experiences which exploit the processes of vision to present you with what is actually there around you are well-placed to serve as an exceptionally basic and reliable way of coming to know about your environment.

On the face of it, attention plays an important part in this story. You come to know about the things you see by attending to them. By attending to the things you see, you come to believe that those things are just as they are represented or presented in your visual experience. You incorporate what figures in your visual experience into your wider cognitive life.

However, work in the cognitive sciences generates a challenge to this intuitive way of thinking about visual experience, because it shows that attention also plays a quite different role. Attention alters visual experience, making it dynamic and variable in ways which do not track real changes or variations in the visible environment. The facts here are not immediately obvious, and many of them are revealed only by controlled empirical studies which I'll explain later. But we can get an intuitive grip on the problem just by considering Figure 1. Focus your attention first on the triangle and then on the square. If you hold the page at the right distance, you can do this all the while seeing every part of the figure. No part of the figure itself changes, of course. Yet your visual experience of the figure changes.

So we seem to face a dilemma. On the one hand, we could give up the claim that your visual experience is simply a matter of your representing, or being presented with, ways your environment *really is*. That is, we could accept that during at least part of your visual encounter with the figure, your experience represents or presents it as having features which it does not actually have. That would explain how your experience differs, while the figure itself remains the same.

Perhaps this proposal does not immediately present itself as plausible, just through reflection on the case of attending to Figure 1. It is not obvious what the error or misrepresentation could be. But we will see that this is the preferred interpretation of some cognitive scientists, with respect to their experimental data about the effects of attention. They argue, for example, that attention to a gap in a shape makes you experience it as larger than it really is (Gobell and Carrasco 2005), and that attention to a colour makes you experience it as more saturated than it really is (Fuller and Carrasco 2006). The effects of attention on visual experience are routine, not rare or out of the ordinary. So if we choose this horn of the dilemma, we give up the claim that visual experience is an exceptionally basic and reliable way of coming to know about your environment, in virtue of representing or presenting things as they actually are.

On the other hand, we could give up the claim that visual experience is simply a matter of representing or being presented with ways your environment *might be*. That is, we could accept that there is more to the qualitative character of your visual experience than simply representing, or being presented with, apparent features of your

environment. That too would explain how your experience could differ, while the figure itself remains the same.

If we choose this horn of the dilemma, we cannot identify visual experience with the sheer gathering of information from your environment. We might still hope to explain visual experience in terms of this gathering of information, in some less ambitious way. But here our prospects are dimmer than we might have thought: there is an aspect of the character of visual experience which does not even co-vary with the information gathered. Furthermore, even this horn of the dilemma leaves us with a challenge when it comes to the epistemic role of visual experience. The intuitive understanding of visual experience I described is motivated by the fact that, when we reflect on the introspectible character of visual experience, we take everything we find to be an aspect of the environment, an aspect of the environment which is presented or represented in visual experience. If this turns out to be a mistake, how can it be that we do not mistake what is not a presented or represented aspect of the environment for such an aspect of the environment? How can it be that the mistake about the character of visual experience does not engender mistakes about the character of the environment?

To escape the dilemma, or at least to appreciate the full range of options open to us, we need to understand the relationship between attention and a certain sort of indeterminacy in visual experience. Indeed we need to understand – by reference to this sort of indeterminacy – what attention is.

It is tempting to suppose that visual experience is only ever experience of maximally determinate properties, the most exact properties which an object has. Accordingly, to describe your visual experience of Figure 1, we should describe the shapes of the triangle and the square, the orientations of the lines and so on, all in maximal detail. This is a mistake – demonstrably so, given some principles I will defend, connecting the contents of visual experience with data about the resolution of visual processes. In fact, visual experience is experience of determinable properties: properties which admit of more specific determinations, as being red is determined by being scarlet, being rectangular is determined by being square, and being 90 to 110 feet long is determined by being 100 feet long. Which determinable properties you experience visually depends in interesting, sometimes surprising ways, on where your attention is directed. Indeed, I will argue that your attending to one thing rather than

another sometimes consists in facts about the determinacy of the properties you experience visually.

By appreciating this sort of visual indeterminacy we can, in principle, see a way out of our dilemma. Suppose we identify the changes in visual experience, which shifts of attention involve, with changes in which properties your visual experience represents or presents you with. If these properties are determinable, rather than maximally determinate, it may be that all the properties you experience visually are actually instantiated by your environment, even though they change while the environment does not. One determinate property determines many determinable properties. So an object with one determinate shape or colour, for example, may take on different appearances, in different experiences of different determinable shapes or colours which the object actually has. The fact that attention makes visual experience dynamic and variable is consistent with the idea that visual experience is a way of being simply open to your environment. This is because the dynamicity and variability of visual experience occurs in an encounter with the environment which is partial or indeterminate, in the sense just introduced.

I will argue that philosophers and scientists have failed adequately to appreciate this sort of visual indeterminacy. I will apply the schematic analysis just described to the experimental data about visual experience and attention which I mentioned above, arguing on empirical grounds that the data are in fact well analysed in this way. Contrary to the experimenters' own interpretations, we need not accept that attention makes visual experience misrepresent shapes, colours and so on. Nor does the solution here require us to accept that there are aspects of visual experience over and above its representing or presenting ways your environment might be. With this analysis in hand, and independently motivated, we can better appreciate the prospects for theories which turn on the idea that visual experience is a way of being simply open to the environment. At least, we can see that the experimental data about attention are consistent with the idea that visual experience is a distinctively basic and reliable source of knowledge about properties of our environment.

However, there are limits to this way out of the dilemma I described. One advantage of my way of arguing for the indeterminacy of visual experience is that it clearly connects this phenomenon with data about the resolution of vision, and with data about the role of visual

resolution in explaining our wider cognitive capacities. There are further experiments (Yeshurun et al. 2008) which suggest that the distinctive phenomenology of visual attention is not exhausted by variations in the determinable properties which visual experience represents. The data here resist explanation in terms of the idea that visual experience is exclusively a way of being open to the environment. On the contrary, they suggest qualifications of that idea, along the lines of the second horn of our dilemma: some aspects of the experience of seeing, as it is modulated by attention, consist in or are explained by further aspects of the role of vision in our conscious lives.

At the same time, we can exploit the evidence here to argue for some specific proposals about how to understand those aspects of attentive visual experience which do not consist in simply representing or being presented with the environment. To put it very roughly for now, these are cognitive, rather than visual, aspects of the overall experience you enjoy when you engage in visual attention. I will argue that, when we understand attentive visual experience in this way, we can understand why these aspects of experience are not misleading with respect to the character of the environment – why they are consistent with the claim that attention and visual experience are an exceptionally basic and reliable source of knowledge.

It has traditionally been assumed that questions about the metaphysics of visual experience must be settled just through *a priori* reflection on philosophical theories about it. I argue that we can exploit the experimental data to understand what form our philosophical theories here should take. By understanding the facts about attention and the indeterminacy of visual experience, we can better understand both the scope and the limits of the intuitive idea that visual experience is a way of being simply open to the environment around us. In turn, we can better understand the nature of visual experience, and we can better understand how that experience forms a distinctive source of knowledge about properties of our environment. This is in part, though by no means exclusively, because we can better understand how work in the cognitive sciences bears on these traditional philosophical topics. What follows in this introductory chapter is a preliminary fleshing out of these claims, together with some remarks on the methodology by which I will argue for them, and a plan of action.

What is attention?

In the scientific literature, the term 'attention' is used in various ways. For example, it is sometimes used to refer to any selective computational or neurological process. Here I am concerned specifically with *conscious* attention – with the kind of attention which plays a familiar role in our introspectible mental lives, and so also plays a familiar role in our folk-psychological understanding of ourselves and other persons. As a rough first pass, we can say that conscious attention is a way of selecting something as a focus of conscious awareness. We can also appeal to some paradigm cases: looking at ϕ, listening to ϕ, noticing ϕ, inspecting ϕ and considering ϕ all entail attending to ϕ (White 1964).[1]

Attending in these ways sometimes involves quite subtle shifts of awareness. Consider Figure 1 again. Direct your gaze at the dot at the centre of the figure, so that the dot is at the centre of your visual field. Even while maintaining fixation on the dot in this way, you can shift your attention first to the triangle and then to the square. Psychologists distinguish between *overt* attention, which in the visual case involves movements of the head or eye, and *covert* attention, which involves no such observable behaviour. Shifts of covert attention have well-tested behavioural consequences (Posner 1980), and they alter your visual experience. The experience of the figure when you attend covertly to the triangle differs in character from the experience of the figure when you attend covertly to the square.

In discussions of attention, it's almost *de rigeur* to cite William James's famous definition:

> Every one knows what attention is. It is the taking possession by the mind, in clear and vivid form, of one out of what seem several simultaneously possible objects or trains of thought. Focalization, concentration, of consciousness are of its essence. It implies withdrawal from some things in order to deal effectively with others.
>
> James 1890/1950: 403

[1] I will often use the expressions 'attend to ϕ', 'attention to ϕ' and other cognates. 'ϕ' should be understood as ranging over an ontologically various domain: the objects of conscious attention include propositions, properties and particulars of various kinds. On the face of it, ϕ need not even be a real – as opposed to merely intentional – object of experience. Plausibly, you can attend in thought to the conspiracy against you, whether the conspiracy is imagined or real.

However, most of those who quote James go on to work with analyses which say something far more specific about what attention is. Recent philosophical work tends to portray attention as essentially *amodal*, or requiring post-perceptual awareness; *disembodied*, or independent of bodily movement; and *agential*, or involving action on the part of the attending subject. Together, these commitments amount to a form of intellectualism about attention. Attention is treated as if it were a form of Cartesian reflection, a deliberate operation of the mind independent of the body and perceptual senses. According to this contemporary picture, attention plays its part only once the senses and the body have played theirs. With the deliverances of the senses all in, the agent exploits attention to choose what to do with them.

I argue that this is a mistake. None of these features is in fact essential to attention. Attention is in some instances a distinctively visual phenomenon; in some instances it requires bodily movement; in some instances it is a passive conscious process, involving no intentional action on the part of the attending subject. Furthermore, we can respect these facts about attention if we work with James's definition. James's definition is successful partly because it does justice to both the commitments and the modesty of our everyday understanding of attention. While it is sufficiently informative to capture the pre-scientific notion, it leaves a certain amount open for empirical and theoretical discovery. We are told that attention consists in a 'focalization of consciousness', but the definition leaves it open what forms this focalization may take, and what mechanisms make it possible.

In the mid-Twentieth Century, psychologists hoped to discover a unique physiological process of attention. By and large, they believed in a distinctive, unitary 'bottleneck' process, by which some of the information registered by the senses is selected for further processing in limited-capacity mechanisms, while other information is discarded (Broadbent 1958; see Figure 2 overleaf). More recently, it has become clear that a bewildering variety of physiological mechanisms realise the selective phenomena usually studied under the heading of attention (Allport 1993). I argue that this holds for the specific case of conscious attention: there is no unitary physiological process of conscious attention; rather, unity is to be found at the level of conscious awareness, in the selective focalization of consciousness which James described.

In this respect, James's approach to attention was somewhat prescient. He went on to offer a theory of the 'intimate nature of the

attentive process' – a theory of the 'physiological processes' which realise conscious attention. He identified both 'anticipatory preparation from within the ideational centres' and 'accommodation or adjustment of the sensory organs' as processes of attention (1890/1950: 411). I will defend James's initial definition independently of this further physiological theory, which is of course rather dated. But the physiological theory has a structural feature which I will argue we should preserve. James recognised that the processes realizing conscious attention are diverse: on the one hand, we may attend in virtue of distinctively post-perceptual, cognitive processes; on the other hand, we may also attend in virtue of changes in the specifically perceptual processes through which we first sense our environment. Correlatively, conscious attention may take various forms. It may take the form of changes and variations in distinctively post-perceptual, cognitive awareness; it may also take the form of changes and variations in the content and character of perceptual experience. Attention, as James defined it, consists in these higher-order properties of first-order episodes of conscious awareness. As it turns out, these first-order episodes may be episodes of conscious cognition, conscious perception, or both, depending on the particular case.

Figure 2

The mid-Twentieth Century view of attention

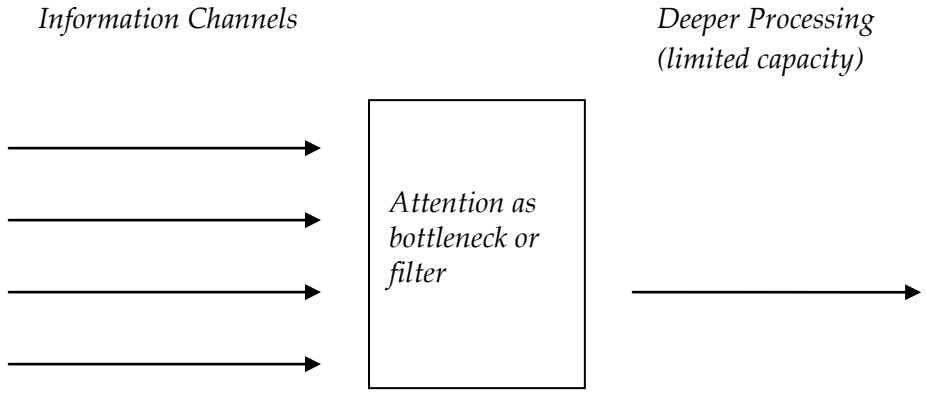

In cases of attentive visual experience, the focalization of consciousness sometimes consists in facts about the determinacy of the properties you experience visually. Where this is the case, your attending

to something consists in your seeing properties of it which are more determinate, relative to the properties you experienced it as having at a previous time, or relative to the properties you concurrently experience other things as having. I argue that empirical studies of the resolution of visual information-processing provide a guide to the determinacy of the properties we experience visually. So interpreted, these studies show that visual experience is indeterminate, in the sense that it is experience of determinable properties. They also show, I argue, that conscious attention sometimes takes the form of changes in the determinacy of the properties you experience visually. In that sense, conscious attention is not essentially *amodal*.

Attention here may be either overt or covert. More determinate properties of an object may become apparent to you either because you direct your gaze at it, or because you attend to it covertly. One advantage of my analysis, over other recent analyses, is that it does justice to the intuitive idea that these are both forms of attention. Where attention is overt, it depends on movements of the body: attention is not essentially *disembodied*.

In some instances, conscious attention is an involuntary, passive phenomenon. In these instances, the mechanisms responsible for changes in the determinacy with which you see lie deep in the visual cortex, and they are controlled by 'bottom-up' responses to ecologically significant changes in the environment, rather than by 'top-down' processes in the control of the person attending. Attending is often something we do, but equally it is often something which just happens to us. In this respect, attending is like weeping or falling, not like murdering or manipulating. Attention is not essentially *agential*.

I think this understanding of attention is worth having in its own right. Attention is a central aspect of our conscious lives, and it is worthwhile to understand what it is, what forms it takes, and how it is related to other aspects of consciousness. But this understanding of attention is also useful because it enables us better to understand both the epistemic role and the intrinsic character of visual experience.

Attention and the epistemic role of visual experience

On the face of it, attention plays an important epistemic role. We often explain why someone comes to know about something she sees, or

why she does not come to know about it, by saying that she pays attention to it or that she does not pay attention to it. For example, suppose that every day on your way to work you walk past houses of various shapes and colours. Despite taking this route every day, you come to know about the shapes and colours of only some of the houses. We can explain this selectivity in what you come to know by saying that, for one reason or another, only some of the houses have occupied your attention. Because you have paid sufficient attention only to these select houses, you have formed beliefs about the shapes and colours of these houses and not the others.

One thing we might have in mind, in offering this explanation, is that you never direct your gaze at some of the houses. Because you never direct your gaze at them, you never see their shapes or colours. As a result, you are not in a position to form beliefs about these shapes or colours. But suppose we stipulate that your gaze has often fallen on each of the houses you pass. Still we might explain the selectivity in what you come to know by saying that you have paid sufficient attention to only some of the houses. In saying this we might have in mind *covert* attention: even though your gaze fell on these houses, you did not attend to them adequately to form beliefs about their shapes or colours.

A natural way to spell out this explanation is to say that, even though you have *seen* all the houses, and your visual experience has represented or presented all their shapes and colours, you have not attended to some of the houses adequately to form beliefs about those shapes or colours; by contrast, when you pay sufficient attention to a house, you come to believe that this house has the shape and colour which your visual experience represents or presents it as having. I think this is broadly the right way to spell out our explanation. Explicitly or implicitly, we often in fact explain what someone does and does not come to know in this way. And this is a plausible explanation in light of both *a priori* and empirical considerations about knowledge and visual experience.

Understood in this way, attention is a distinctive source of knowledge, because attention is a means of fixing beliefs which inherit their content from visual experience: through attention, you come to believe that things have the very properties which are represented or presented in visual experience. You believe that the house has a certain shape and colour, say, because that shape and colour figure in your visual experience of it. Beliefs of this kind have a distinctive epistemic status. They inherit the distinctive reliability of visual processing (Dretske 1997),

or the distinctive justification or warrant provided by conscious vision (McDowell 1994; Burge 2003). This role for attention has figured surprisingly little in discussions of the epistemic role of visual experience. But once you think about it, it's clear that an understanding of the role of attention is essential to a full understanding of the epistemic role of visual experience.

As I said, we'll see that attention in fact plays a dual role here. Covert attention dictates not only what you take up in thought and belief, but also the character and content of visual experience itself. There is some temptation to assume that covert attention does not play this second role – to assume that visual experience consists in the presentation of a stable visual array, an array over which attention roams without altering the array itself. And there is a danger that this assumption gets presupposed in our philosophical theories about visual experience and visual knowledge.

For example, consider MGF Martin's analysis:

> The notion of a visual experience of a white picket fence is that of a situation being indiscriminable through reflection from a veridical visual perception of a white picket fence as what it is.
>
> Martin 2006: 363.

Martin intends this style of analysis to apply to visual experiences in general: to have a visual experience as of a phenomenon ϕ is to be in a situation which is indiscriminable through reflection from a veridical visual perception of ϕ as what it is. 'Reflection' consists in attending to your experience, and so in attending at least covertly to the objects of that experience (cf. Martin 1998). So on the face of it, the analysis presupposes that covert attention does not affect what you experience visually. To see this, suppose that covert attention does affect what you experience visually: when you attend to something you experience ϕ, and when you don't attend to it you experience χ. Martin's analysis seems to rule this out, treating both the experience of ϕ and the experience of χ as experiences of ϕ. For reflection on your visual experience of χ will alter it into a visual experience of ϕ, hence into an experience indiscriminable from a veridical visual perception of ϕ as what it is.

Similarly, work in epistemology which takes attention seriously tends to focus only on attention's role downstream of visual experience, so

to speak, ignoring the role it plays in fixing the contents of visual experience. Imogen Dickie's (2010) discussion is a helpfully clear example. Dickie argues that visual experience furnishes a uniquely basic kind of knowledge of physical objects:

> A subject, S, is 'acquainted' with an object, *o*, iff S is in a position to think about *o* in virtue of a perceptual link with *o* and without the use of any conceptual or descriptive intermediary.
>
> Dickie 2010: 213

I don't dispute Dickie's account of perceptual acquaintance. But her argument turns on the claim that 'pre-attentive processing delivers up a field … over which your attention roams'; 'pre-attentive processing is already serving up a scene parsed into objects that we can then attend to' (219-220). The role of attention is just to make selected parts of this experienced scene, including the physical objects which visual experience represents, available as the subject-matter of thought and belief:

> Consider your visual field as it is laid out in front of you when you look at this page. Your visual field contains much more information than you can access for conceptual thought. For example, you can see a large number of words on the page. But at each moment you will be able to tell what only a few of them are. 'Selective attention' is the process by which a specific part of the visual field is highlighted as a part accessible to conceptual thought.
>
> Dickie 2010: 216

In this respect, Dickie's account sits well with the 'bottleneck' view of the mechanisms of attention (Figure 3).

Now attention does make parts of a scene available as the subject-matter of thought and belief. So I cannot object, in principle, if someone defines a notion of attention in terms of just this role. But in fact the form of attention which is familiar in our conscious lives also plays a further role, in fixing the contents of visual experience. With respect to this familiar form of attention, there is no pre-attentively fixed field over which attention roams. This is obviously the case with overt attention: by directing your gaze at something, you can make it an object of visual experience. But we'll see that it is also the case with covert attention: even

once the facts about where your gaze is directed are settled, covert attention has further effects on what you experience visually.[2]

Figure 3

Visual experience and the 'bottleneck' view of attention

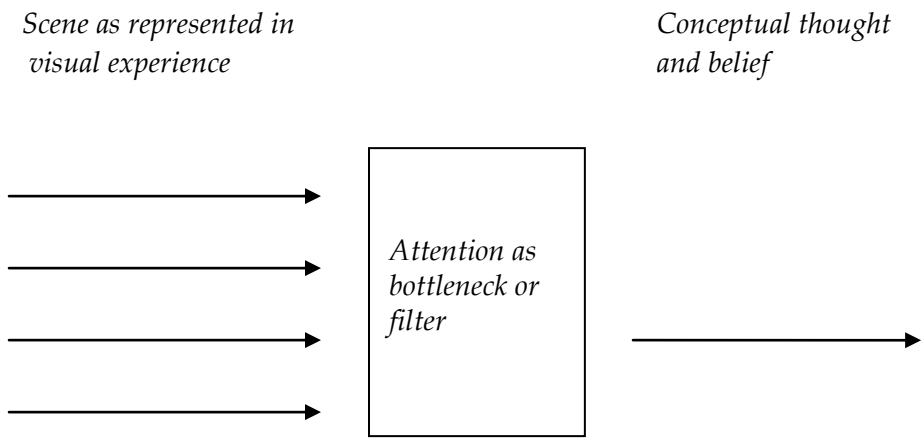

Once we acknowledge that attention plays this dual role, we face a challenge concerning the epistemic role of visual experience. The challenge lies in the fact that attention makes the contents of visual experience vary in ways which do not track variation in the visible environment.

Some variations in visual experience are less problematic than others. For example, if you direct your attention to part of a scene, you may see things which you did not see before. This generates no great challenge concerning the epistemic role of visual experience: you are now in a position to come to know about parts of the scene which you were not in a position to come to know about before.

Other variations in visual experience, though important and challenging, have been discussed in detail by others. For example, depending on the lighting conditions and on the angle from which you look at something, the colour and shape of the thing take on different

[2] Most relevantly to Dickie's discussion, conscious spatial attention is one factor in determining which physical objects you experience visually (Campbell 2007). However, my focus will instead be on the role of attention in determining which properties of things we experience visually.

appearances in your visual experience. These variations are consistent with the idea that visual experience presents a scene just as it is, and so puts you in a position to know about the scene just on the basis of your experience. The lighting and the angle are both aspects of the scene perceived. If we recognise that these varying aspects of the scene are perceived along with the invariant colour and shape, we can accommodate the idea that visual experience presents things just as they are, even while visual experience varies in this way (Noë 2004).

I will focus on variations in visual experience which are distinctive, in that they occur in different visual experiences of a single part of a scene, and while nothing of which you are aware in that part of the scene actually varies. These are variations in the appearance of seemingly simple properties of the things you see, such as their shapes and colours.

James claimed that attention alters visual experience, highlighting attended properties by giving them a 'more intense' appearance (1890/1950: 425). In recent experiments, Marisa Carrasco and her colleagues have confirmed this claim. They found, for example, that cuing attention to a colour makes it appear more saturated than it really is (Fuller and Carrasco 2006), and that cuing attention to a gap in a shape makes the gap appear larger than it really is (Gobell and Carrasco 2005). Yet James remarked that 'the intensification … never seems to lead the judgement astray' (426): even given these effects of attention on visual experience, visual experience seems to be a reliable source of knowledge. The challenge is to understand how this could be. How could attention so alter visual experience without distorting the way visible properties show up in it, and making visual experience a systematically misleading source of belief? How could attention and visual experience form a reliable source of knowledge?

Now under demanding experimental conditions, Carrasco's subjects do make false judgements about the properties of stimuli to which their attention is cued. This, combined with controls designed to rule out the possibility that subjects' mistakes are merely a product of post-perceptual misjudgement, is her evidence that attention induces variations in visual experience. And her conclusion is pessimistic, at least from an epistemologist's point of view: she concludes that attention produces 'nonveridical percepts … by emphasizing relevant details at the expense of a faithful representation' (Carrasco et al. 2008: 1162). If this were right, then someone who believes that things are just as her visual

experience presents them would form systematically false beliefs. Visual experience and attention would not be a reliable source of knowledge.

I argue that we can resist this pessimistic conclusion, and do so in an empirically respectable way, if we reflect carefully on the nature of the properties we experience visually, and on the connection between visual experience and the varying resolution of visual processes. As I said, I will argue that visual experience represents or presents us with determinable properties. Because visual experience is indeterminate in this sense, attention may change the way things appear visually, without making experience 'nonveridical'.

I exploit the notion of a property space, and the various transformations which are possible within it, to show how to interpret Carrasco's data and James's claim in these terms. Cuing attention to colours and shapes gives them an appearance which other, more intense colours and shapes also have when unattended. In that sense, attention gives these properties a more intense appearance, but none of these appearances need be nonveridical. Furthermore, we can explain in these terms why attention leads subjects to make mistakes in Carrasco's demanding experimental tasks, even though – as James said – attention 'never seems to lead the judgement astray' under normal circumstances. The demanding task forces subjects to go beyond what their experiences present them with, in reaching judgements about the stimuli they see. By contrast, someone who simply believes that things are as her experience presents them will not fall into error, even where her attention is cued. I argue that this account is as empirically plausible as an account in which attention makes visual experience illusory. The empirical evidence is consistent with the idea that visual attention is a reliable source of knowledge.

There is also a broader issue here, about the relationship between experiment, judgement and visual experience. Carrasco's work is distinctive partly because she connects experiments in psychophysics with claims about the contents of visual experience. The psychophysical data consist in subjects' discriminations between visible stimuli, where discrimination is a form of judgement. Psychophysicists use these data to investigate processes in the visual system. But it's unobvious how we should connect their work with claims about the content of visual experience. I argue that, to do this in a principled way, we need to understand the role of determinable properties in visual experience and visual discrimination.

My interpretation of Carrasco's experiments does not require that we ascribe to visual experience any features over and above its representing, or presenting you with, aspects of your environment. (Indeed, I argue that my interpretation is more successful than Ned Block's (2010) interpretation, which does turn on ascribing such features to visual experience.) This might suggest that, in light of the empirical evidence, we can fully vindicate the intuitive idea that visual experience is a way of being simply open to your environment. However, I argue that once we take full account of the evidence about attention and the resolution of visual processes, this will not quite do. The distinctive phenomenology or conscious character of attending to the things you see is not well correlated just with variations in which determinable properties you experience visually. To appreciate this, we need to understand how attention is connected with visual experience, and what attention itself is.

Attention and the nature of visual experience

Seeing is a way of being conscious. It has what is variously called a qualitative, subjective or phenomenal character, or a phenomenology; there is 'something it's like' to see.[3] Famously, this makes it hard to understand how episodes of seeing could be realised in the brain, and more generally how they could form part of the world described by the natural sciences.

To understand this phenomenon of conscious visual experience, we need to understand attention. Over the last decade or so, empirical studies of attention have inspired various motivations for this idea, each of them controversial. Some theorists argue that attention is essential to conscious

[3] It is sometimes said that 'what it's like' is a foolish expression in this context, because it does not mean what philosophers here take it to mean. But presumably philosophers do not think the expression *ordinarily* means what they use it to refer to. Rather they use it to advert, in a quasi-technical vocabulary, to a phenomenon of which they hope we all have a notion, at least once we reflect carefully on our mental lives. In part, the expression 'what it's like' is a way of alluding to Nagel's (1974) canonical discussion. That said, I think we should take care to use the expression in as thin a sense as possible, without prejudice as to the nature of consciousness. For instance, in using the expression I do not imply that there are features of consciousness which, necessarily, are appreciable only from the first-person perspective.

vision – that to experience something visually is, in part, to attend to it (Mack and Rock 1998). Others argue that attention is necessary and sufficient for consciousness (Prinz 2003), or that conscious experience is realised in the brain by the mechanisms of attention (Posner 1998). And others argue on methodological grounds: we can operationalize attention, but not conscious experience, so a science of attention is the closest we can come to a science of consciousness (Deheane and Changeux 2004).

However, there is a more obvious reason why we need to understand attention, if we're to understand conscious experience. The phenomenology of attention makes a ubiquitous and fundamental contribution to the phenomenology of human conscious experience. James noted the 'perpetual presence of selective attention', and remarked that 'without it the consciousness of every creature would be a gray chaotic indiscriminateness, impossible for us even to conceive' (1890/1950: 401-2). Perhaps there could be a consciousness which was not qualified by attention, a consciousness for which nothing was focal and nothing peripheral. Philosophers have sometimes encouraged this idea, by attempting to characterise our actual conscious experience without mentioning attention. But a consciousness of this kind would in fact be quite different from our own.

To illustrate, first concentrate your attention on the triangle in Figure 1. Then concentrate instead on the square. Then listen carefully to any sounds you can hear at the moment. Then divide 216 by 18, or call to mind the facial features of Barack Obama. In every case you are conscious of the triangle, but in every case your experience of the triangle differs quite fundamentally. One object of consciousness is selected at the expense of others, as a focus of your conscious awareness. This selection may be intra-modal, in the sense that one part of the figure is selected at the expense of others within the single modality of visual experience. Or the selection may be inter-modal, for example occurring among objects of visual and auditory experience. Or it may occur among the various objects of conscious perception and conscious cognition: your focus may shift between things which are present to your perceptual senses and things which are not. As James noted, this selective structure is pervasive in our conscious experiences. At the very least, structuring around a focus of attention is a fairly basic aspect of the qualitative character of many ordinary conscious experiences.

This gives us a strong motivation for thinking that, to understand conscious experience, we need to understand attention: the

phenomenology of attention makes a ubiquitous, fundamental contribution to the qualitative character of conscious experience. I will argue that by understanding attention properly, we can achieve a better understanding of the nature of visual experience. In particular, we can better understand what has been called the *transparency* of visual experience, together with consequent issues about the metaphysics and physical realisation of consciousness.

To put it roughly just for now, your visual experience is transparent to the extent to which the following is true: when you reflect introspectively on your visual experience, what's revealed to you is exhausted by the particulars and properties which form the objects of that experience. The thought that visual experience is thoroughly transparent has suggested to philosophers that visual experience is a way of being simply open to the visible environment in the following sense: visual experience consists entirely in representing or being presented with particulars and properties of the visible environment. Understanding the qualitative character of visual experience in this way holds the promise that we might explain this qualitative character in terms of the information-gathering carried out by visual processes.

In some theories, the explanation is that the qualitative character of visual experience is *identical* with the gathering of information by some visual processes. Visual experience is identified with representation of the visible environment (Tye 1995), and this representation is identified, in turn, with some way of bearing information about aspects of the environment. For example it's argued that for a neural state to represent ϕ is for that state to depend asymmetrically on ϕ (Fodor 1990), or that for a neural state to represent ϕ is for that state to have the biological function of indicating ϕ (Dretske 1995).

In other theories the explanation is less ambitious. For example John Campbell (2002) argues that visual experience consists in a 'simple relation' to physical objects and their properties. The relation is simple in that it admits of no further analysis. For example, it cannot be analysed in terms of representation, in the way in which some theories analyse the relation of seeing a physical object. Campbell claims that visual information-processes carry out 'a complex adjustment that the brain has to undergo, in each context, in order that you can be visually related to the things around you' (Campbell 2002: 119). This approach does not identify visual experience with visual information-gathering. But to the extent to which information-gathering processes explain which visual relations you

stand in, we have an explanation of the qualitative character of visual experience: in virtue of the fact that vision gathers information about certain aspects of your environment, you are visually related to those aspects of it; to enjoy visual experience is just to be so related – 'the qualitative character of the experience is constituted by the qualitative character of the scene perceived' (2002: 115).[4]

James's comment about the 'perpetual presence of selective attention' raises a serious challenge for these theories of visual experience. When you reflect introspectively on an episode of visual experience, there are (at least in many ordinary cases) facts about the qualitative character of the experience over and above the facts about which aspects of your environment the experience presents or represents; there are, in addition, the facts about how focal or peripheral the various aspects of your environment are in your conscious experience.

The idea that attention raises a challenge for these theories of visual experience is not new (Chalmers 2004; Speaks 2010). But I think it has been a bar to progress here that philosophers tend to work, implicitly or explicitly, with a false analysis of attention. In particular, it's been a bar to progress that the connections between attentive experience and the resolution of visual processes have not been adequately appreciated.

One response to the challenge I have described is to try to capture the distinctive phenomenology of attention in terms of variations in the determinacy of the properties experienced (Nanay 2010). If the degree to which something is focal to visual experience co-varies in every case with the determinacy of the properties you experience it as having, then we might capture the phenomenology of visual attention in terms of variations in the properties which visual experience presents or represents. So we might defend the claim that visual experience consists entirely in representing or being presented with particulars and properties of the visible environment.

I will argue that this hypothesis about attention is subject to an empirical test – a test which has in fact been carried out, though with a somewhat different aim. Yaffa Yeshurun and her colleagues developed a way of testing the effects of attention on the spatial resolution of visual processes (Yeshurun and Carrasco 1998; Yeshurun et al. 2008). The test

[4] Representational theories of visual experience could likewise appeal to this more modest form of explanation; they need not identify conscious visual representation with information-gathering by visual processes (e.g. Searle 1989).

shows that the performance-enhancing effects of attention, which (I argue) co-vary with the distinctive phenomenology of attention, cannot be explained just in terms of the idea that attention increases visual resolution. The distinctive phenomenology of visual attention does not co-vary just with changes in the determinacy of the properties you experience visually.

In light of this, I argue that there are facts about the conscious character of episodes of visual experience over and above the facts about which aspects of the environment they present or represent. There are, in addition, facts about how focal things are in your conscious experience – facts about the degree to which you are conscious of things, rather than about what you are conscious of.

However, I exploit the Jamesian analysis of attention sketched above, to argue that philosophers have misconstrued the challenge which attention poses here. The phenomenology of attention poses no *distinctive* challenge to the claim that visual experience is simple openness to the environment, since attention is not a distinct form of consciousness, over and above conscious perception and conscious cognition. Rather, the challenge here is just to accommodate an obvious fact: that conscious cognition of ϕ and visual experience of ϕ are compossible. I conclude by exploring some of the options for understanding this fact about consciousness, and its consequences for the project of explaining visual experience in terms of an information-processing story about the mind.

Methodology

The phenomena I investigate in this dissertation are aspects of consciousness. They play familiar roles in our introspectible mental lives, and so also play familiar roles in our folk-psychological understanding of ourselves and other persons. I pursue questions both about the nature of these phenomena, and about their role in epistemology. The method by which I pursue these questions is distinctive partly because of the way in which I connect them with empirical evidence about the processes of vision.

Presumably, there are many questions in epistemology and the philosophy of mind which cannot be settled just by findings in the cognitive sciences, or indeed by any empirical test. But in general, philosophical theories about conscious, personal level psychology are

responsible to our best scientific understanding of the mind. We should accept this independently of the attitude we take to questions about the ontological relationship between states and episodes of personal-level psychology on the one hand, and on the other hand the neural and information-processing phenomena typically investigated by cognitive scientists. Perhaps this is obvious, but it bears emphasis nonetheless. By sociological accident, many philosophers who draw on scientific results in their understanding of the mind take it that personal-level and neural phenomena will turn out, in the final analysis, to be identical. But those who deny this or find it incomprehensible do mostly accept that personal-level psychology depends in some way on these lower-level phenomena. What we say about the personal level must respect this dependence, if it's to have a hope of being true.[5]

This provides one sort of justification for my appeal to empirical results in pursuit of philosophical questions. For example, I argue that what we say about the properties we experience visually must be constrained by what we know about the resolution of visual information-processes: we should not say that we are visually conscious of properties more determinate than visual information-processes – with their limited and varying resolution – could make possible. This line of argument does not turn on any identification of visual experience with visual information-processes. The line of argument is independent of any such reductive theory.

There is also a further, more direct way in which I think empirical results bear on the claims we can legitimately make about the contents of visual experience. Vision scientists have devised experimental controls which distinguish between two kinds of effects: (i) effects on the signal which the visual system processes in response to a stimulus; (ii) effects on the decision-making processes which mediate between the visual signal and observable behaviour. I argue that we can exploit these controls to a somewhat different end. Where effects on behaviour are found to be the upshot of effects on the visual signal, and where vision is conscious, we have reason to believe that we're dealing with effects on visual experience, rather than just effects on the judgements we make in response to our

[5] This is far from being a novel thought. For example, Aristotle maintained that you cannot understand seeing properly without understanding the visual system; he also argued that what it is to see cannot be captured or defined in terms of the visual system – just as the house-builder must understand bricks, even though houses cannot be defined in terms of bricks (*De Anima* I.1; *Physics* II.2).

visual experiences. For example we can understand where, and in what ways, attention alters the contents of visual experience.

However, the appropriate attitude here is not just a passive acceptance of the scientists' interpretations of their results. The appropriate response is a dialogue with cognitive science. So for example where the scientific results are not, on the face of it, consistent with the view of the mind with which we usually work in philosophy (as in the case of Carrasco's experiments), the onus is on philosophers to take the empirical results seriously. But equally we may respond by questioning the scientists' assumptions about how their data bear on conscious visual experience. For example, I resist Carrasco's interpretation of her results by showing that she does not take full account of the role of determinable properties in visual experience and visual discrimination.

Indeed, when vision scientists distinguish effects on the visual signal from effects on decision-making, they usually argue by reference to *signal detection theory*. This theory models visual judgements and discriminations as the outputs of a statistical decision process which takes visual signals as its input. I argue that, to understand what this picture of the mind tells us about visual experience and personal-level discrimination, we need to appreciate the role of determinable properties in visual experience. By doing so, we can see that the scientists' results comport much better with the philosophers' view of the mind than we might at first glance suppose.

I have already said that I will prescind from one important debate in the philosophy of mind: the debate about the ontological relationship between conscious episodes and the phenomena typically studied by cognitive scientists. The same goes for two other important debates in the philosophy of perception. One is the debate – mentioned above – about whether episodes of conscious perception are representational, or instead consist in simple relations to the environment. There is of course much of interest to be said about this in connection with the issues I do discuss. But the debate is somewhat intractable, and I don't want what I say here to be hostage to one side of it or the other. To that end, in Chapter 2 I sketch a way of discussing the contents of visual experience without taking a stand on this issue.

A second debate concerns the sorts of properties which may figure in the contents of visual experience. Most philosophers accept that colours and spatial properties are among them, but it's controversial whether natural kind properties, for example, are too. To avoid this debate, and to

keep things fairly constrained, I will focus on colours and spatial properties.

Also in order to keep things fairly constrained, I will focus on the relationship between attention and the visual experience of properties, rather than physical objects. There is much of interest to be said about the relationship between attention and the visual experience of objects, but lots of it has already been said (esp. Campbell 2002), and the combined project would get out of hand. Finally, some of what I say here could also be applied to conscious experiences in the non-visual perceptual modalities. But I leave these connections for another day, because here too the combined project would be enormous.

Very recently, and late in my work on this dissertation, there has been a small flourishing of philosophical work about attention (see especially Mole 2011, and the papers in Mole, Smithies and Wu 2011). In what follows I criticise some aspects of this work, and I use it as a foil where it's helpful to do so. But I don't pretend even approximately to do justice to the recent theories I mention, or to address all the complex discussions used to motivate them. That would require far longer than I have available, both in time and in text.

The Plan

In Chapter 2, I argue that visual experience is indeterminate in the sense I have introduced, drawing on data about the limited and varying resolution of visual processes. This material provides a crucial premise in most of the subsequent arguments of the dissertation.

I then turn to the nature of attention, and to its connection with visual experience. In Chapter 3, I criticise some previous discussions of the relationship between attention and visual experience, in both philosophy and cognitive science. I define a notion of constitutive connection among psychological states, and assess the arguments others have proposed about the question whether visual experience is constitutively connected with attention. I argue that the arguments in the existing literature make no real progress with this question.

In Chapter 4, I defend James's definition of attention, and criticise some alternative analyses which philosophers have offered recently. I argue that, once we understand the nature of attention and its connection

with visual experience, we can answer the question of Chapter 3: in some instances, visual experience is connected constitutively with attention.

In Chapter 5, I examine James's claim that attention leads to '*more intense*' visual appearances, together with Carrasco's studies confirming this claim. I show how James's observations and Carrasco's results do not undermine the epistemic role of attention and visual experience. The psychophysical results are consistent with the idea that attention and visual experience form a distinctive source of knowledge.

Finally, in Chapter 6, I show how the evidence about attention and visual resolution bears on the nature of visual experience: in light of this evidence, we need to qualify the intuitive idea that visual experience is a way of being simply open to the environment.

Chapter 2

VISUAL INDETERMINACY

We can distinguish between a determinable property and its determinations. The property of being octagonal determines the property of being shaped; being crimson and being scarlet each determine being red, being red determines being coloured, and being six feet tall determines being more than five feet tall. Which properties figure in human visual experiences? Is awareness of some properties distinctive or definitive of visual experience?

There is a traditional and persisting assumption about this, which I will call *Determinacy*:

> *Visual experience is never experience of determinable properties, without also being experience of determinations of those properties.*

The idea is that visual experience is always experience of maximally determinate properties. Hume famously wrote,

> 'tis confest, that no object can appear to the senses; or in other words, that no impression can become present to the mind, without being determin'd in its degrees both of quantity and quality. The confusion, in which impressions are sometimes involv'd, proceeds only from their faintness and unsteadiness, not from any capacity in the mind to receive any impression, which in its real existence has no particular degree nor proportion.
>
> Hume, 1740/1978, I.1.vii

There are various ways to read the passage, but on one interpretation Hume here commits to *Determinacy*. According to Hume, sense experience consists in impressions being present to the mind. Impressions are particular items, but the theory surrounding them generates a commitment to *Determinacy*, a principle about the visual experience of properties. Like ordinary objects ordinarily conceived, impressions

instantiate determinable properties only where they instantiate determinations of those determinables. Your impression is shaped only if it's shaped in a determinate way – rectangular, for example. And for an impression to have a quality or property, that quality or property must be present to the mind. So if a determinable property is present to you, a determination of that property is also present to you. Restricting to the visual case, we get *Determinacy*.

Although few today would offer this argument for it, *Determinacy* persists as an assumption in the philosophy and science of perception. It is tempting to think that the awareness of maximally determinate properties is part of what's distinctive of visual experience, in contrast with thought or judgement. Suppose I report something's shape by saying, 'I saw that it had between ten and fifteen sides, but I can't be more specific than that.' It's natural to assume that some maximally determinate shape must have been visually apparent to me; the most specific judgement I could make abstracted from that determinacy, leaving me only with the thought that the thing had between ten and fifteen sides. That sort of abstraction, it's tempting to suppose, is the hallmark of mere judgement, where by contrast maximal determinacy is the hallmark of visual experience.

This line of thought is a mistake, or so I will argue in this chapter. The determinacy of the properties which figure in visual experience is deeply variable. Violations of *Determinacy* are the norm. Indeed, with respect to at least some properties, including the spatial properties of length, shape and location, visual experience is *never* experience of maximally determinate properties. At least, we must accept these claims if we are to understand visual experience as a perceptual encounter with instantiations of properties in the environment around us – an ambition which Hume did not share, but which many philosophers and scientists do harbour today. However, before proceeding to my argument against *Determinacy*, I want to borrow an idea from Hume – an idea which he expresses in terms of something's being 'present to the mind'.

2.1 Visual Experience of Properties

For Hume, the qualitative character of experience is given by the qualitative character of what is present to the mind. If we want to describe the qualitative character of your visual experience, we must describe what is present to your mind when you see. We can adopt this idea without

accepting some of Hume's further claims about what is present to the mind. For example, we need not accept that what is present to the mind consists in visual impressions.

On the face of it, it seems that the properties we must describe, in order faithfully to describe our visual experiences, are properties of the environment around us – for example the colours, shapes, and locations of the objects which make up a visible scene (Strawson 1979).[6] I will focus on the idea that properties in this class figure in our visual experiences in a certain distinctive way. Where a property *P* so figures in the experience of a normally-sighted human subject *S*, the following condition is met:

Presence

> A faithful description of the qualitative character of *S*'s visual experience must mention *P*.

As I said in Chapter 1, I will focus on colours and spatial properties, though I do not assume that these are the only properties which meet the condition specified.

That some properties meet this condition is a fairly standard assumption in contemporary philosophical work.[7] Different philosophers talk in different ways about the visual experience of properties: you see things as red; things look red to you; your visual experience represents things as red; the redness of things is manifest to you; you see redness; you see red. These different expressions often encode different theories about the nature of visual experience, and about the role of properties in visual experience. However, philosophers talking in these different ways are often talking about the same thing, though they presuppose different theories about what it fundamentally is. Often – though by no means all of the time – they are talking about the fact that properties meet the condition of *Presence* for a subject.

Consider two examples. John Campbell (2002) argues that visual experience is a 'simple relation' to physical objects and their properties –

---

[6] Hume did not deny this, of course. Rather he argued (from illusion) that the naïve view is mistaken (1740/1978: I.4.ii). I will not rehearse the standard responses to his argument.

[7] For some dissent, see Travis 2004 and Peacocke 1989. I have neither the space nor the resources to address Travis's arguments here. In Chapter 5 I show how to resist Peacocke's argument, which is based on phenomenal *sorites* cases.

the relation of seeing those objects and properties.[8] Broadly speaking, Campbell adopts Hume's idea about the visual presence of impressions, but applies it instead to physical objects. According to this proposal, 'the qualitative character of the experience is constituted by the qualitative character of the scene perceived' (114). So to describe the qualitative character of the experience, we must describe the properties of the experienced scene.

A different proposal is that the relation of seeing should be analysed into a representation of objects and their properties, together with some further (perhaps causal) connection between the representation and objects which instantiate those properties (Searle 1989; Tye 1995). Here too, it is often taken for granted that some properties meet the condition of *Presence*. For example, consider Susanna Siegel's (2006) argument that natural-kind membership is among the properties which visual experience represents, along with the colours and shapes of things. Siegel argues on the grounds that representation of each of these properties makes a difference to the phenomenology or qualitative character of visual experience. And she takes her argument to establish what others have tried to establish by insisting on descriptions of the character of visual experience which mention these properties (485).[9]

In both these theories, the idea that properties meet the condition of *Presence* plays a similar theoretical role: that a property *P* meets this condition for a subject, in the way each theory states, explains how he subject comes to know about *P*. Campbell's claim that visual experience is a relation to the intrinsic properties of things is intended to explain how we come to know about these properties (137-156; 235-252). And Siegel takes it that visual experiences 'provide justification for believing' propositions about the properties which they represent (487).

---

[8] The idea is that in seeing an object you are related to its properties insofar as the object instantiates them: you cannot bear the visual relation to properties which are not instantiated by your environment.

[9] Properties might meet this representationalist version of *Presence*, even if there is more to the qualitative character of visual experience than we can describe just by describing what is represented. Properties might meet this version of *Presence*, whether or not the qualitative character of visual experience consists in, or supervenes upon, what visual experience represents. In Chapter 6 I assess various supervenience theses about visual phenomenology and visual representation, but for now everything I say is neutral with respect to them.

Again in both cases, visual experience of a property *P* is taken to be a distinctive kind of awareness of *P*, in terms of which wider knowledge about *P* may be explained. According to Campbell, to be visually related to a property *P* is to enjoy a kind of awareness more basic than knowing about *P*, in the sense that being visually related to *P* explains and justifies knowledge about *P*, but neither entails nor is entailed by knowing about *P* (122). According to Siegel, your knowing or recognising *P* on the basis of visual experience does not by itself entail that your visual experience represents *P*. Rather, she treats it as a substantive question which of the properties recognised on the basis of visual experience are represented in visual experience itself. So both these theories appeal to a distinction between visual experience of *P*, and knowledge about *P* which is acquired through vision. *Presence* likewise reflects this distinction: it concerns the faithful description of *visual* experience, not the subject's overall experiential state. I'll return to this feature of *Presence* shortly.

The different theories which Campbell and Siegel defend are theories of the same phenomenon: that of a property's meeting the condition of *Presence* for a subject. I think this phenomenon should be central to our understanding of visual experience, and I want to make progress with it without first settling the intractable questions about the nature of visual experience. So when I say that *S* sees or experiences *P*, I will mean – stipulatively – that *P* meets the condition of *Presence* for *S.* Whatever expression we use, it is important to make this stipulation. One might take much more liberal views of the conditions on seeing or experiencing *P*, and likewise of the conditions on being visually related to or visually representing *P*. To ease the exposition of some disputes, I will often say that *S*'s visual experience represents *P*. Here too I will mean just that *P* meets the condition of *Presence* for *S*. I do not assume that visual experience should in fact be analysed in terms of representations, and my positive proposals could readily be translated into a different framework such as Campbell's.

I think these philosophical theories reflect a pre-theoretic commitment: in everyday psychological explanations, we routinely exploit the idea that a property meets the condition of *Presence* for a subject, in explaining what she comes to know. An example will help to bring out some key features of this idea.

Suppose you're playing snooker, and when you come to the table the balls are arranged as in Figure 4. You have to pot the red ball first. How do you know which ball to aim at? There are several things we

might naturally say about this. We might say that you can see which of the balls is red; or that only one of the balls looks red to you; or that you see the red colour of the red ball and the other colours of the other balls. This is a jumble of ideas about our psychology, none of them synonymous with another. So we could do with a way of teasing them apart.

Figure 4

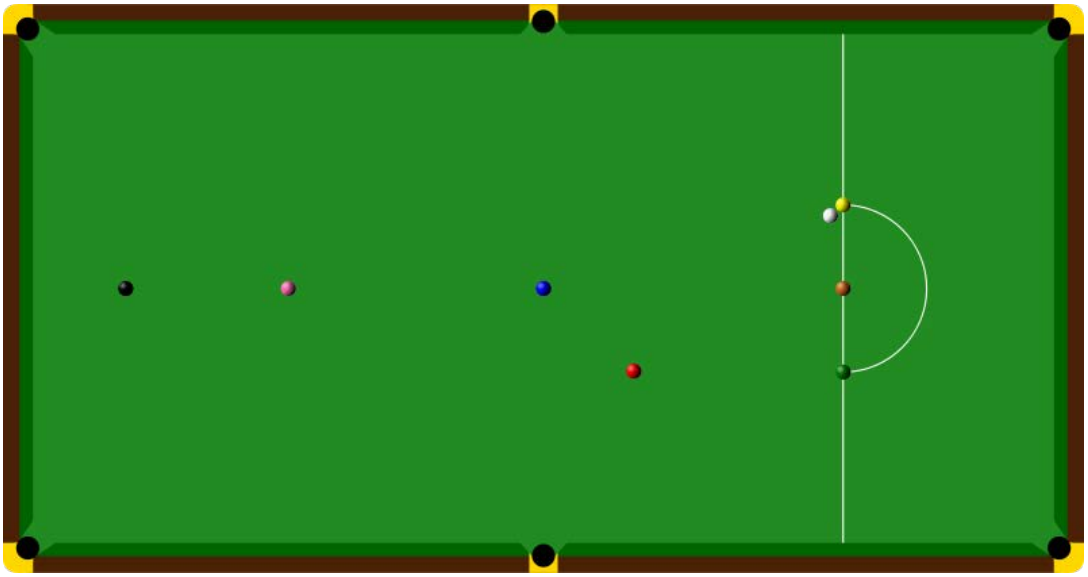

To expose the anatomy of an ordinary explanation, it often helps to think about a case where the explanation breaks down. In a professional snooker match in 2008, former World Champion Peter Ebdon fouled by potting the brown ball when he should have potted a red. Why? Ebdon suffers from Daltonism, the common form of colour-blindness due to retinal insensitivity to variations along the red-green dimension of colour. Most of us know that the brown ball is not red because we see its distinctive brown colour. Ebdon does not see this colour in the way that most of us do. His experience of colour does not distinguish between some shades of brown and red, where most people's experiences would distinguish between them.

How are we to describe this difference between our experiences? Characterising the experiences of the colour blind correctly is a bit of a minefield (see Broackes 2010 for discussion). But I'm more concerned to

draw out some distinctions in folk psychology than I am concerned with the correct theory of Daltonism. And it's natural, I take it, to suppose that we should mention the specific brown colour of the ball in describing the qualitative character of most people's visual experiences, but withhold mention of it in describing the qualitative character of Ebdon's visual experiences. To explain why Ebdon failed to know which ball was brown, where most of us would have known it, we can appeal to the idea that the specific brown colour of the ball meets the condition of *Presence* for most of us, but not for Ebdon.[10]

Now what Ebdon lacks, and most of us have, is not simply visually-based knowledge about the colour of the ball. He does usually know which ball is brown, and he knows this on the basis of his visual experiences. For example, he keeps track of the brown by reference to the path it follows from a starting-position dictated by the rules of snooker. In the 2008 match, Ebdon's way of coming to know broke down. He failed to keep track of the ball's location. For most people, failing to keep track of the ball's location would not have led to a foul. So our explanation of Ebdon's behaviour turns on something he lacks *even when he knows which ball is brown*. The colour of the ball does not figure in his visual experience in the way it would for most of us. Because his visual experience so differs from most people's visual experiences, Ebdon does not know which ball is brown in the way in which most people would. What he lacks is not simply visually-based knowledge, but a certain distinctive kind of visual experience of a property.

When we offer this explanation of the fact that Ebdon mistook the brown for a red, where most of us would not have done, we are appealing to the idea that visual experience makes available to us a general feature or property of the ball, a feature which other things could have. What Ebdon lacks, and most of us have, is not a capacity to see any particular object. He saw all the balls just like the rest of us. You might argue that Ebdon failed to see a particular *property-instance*, namely the instance of that distinctive brown which most of us would have seen. I don't know whether that would be right. Perhaps Ebdon saw the instance of brown;

[10] If what we say is to be true, it may be important that we mention the specific brown colour of the ball, rather than its sheer brownness. Broackes argues that Daltonists do, under some circumstances, see some shades of green, red and brown. Still, there may well be a fairly determinate shade of brown which Ebdon did not on this occasion see, and which most of us would have seen. However, for ease of exposition I will use simply 'brown' to refer to the specific brown colour of the ball which most of us would see.

he just didn't see it *as* an instance of brown. But suppose we accept that Ebdon did not see the instance of brown. The obvious reason for accepting this is just the view that seeing property-instances requires awareness of which properties are instantiated. So on the face of it, even if Ebdon's behaviour is explained by his failure to see a particular property-instance, this failure is at root a failure in his awareness of a general feature or property.[11]

It is worth forestalling a potential confusion here. If you try to understand the idea that we are visually aware of properties on the model provided by our seeing particular objects, the idea can seem absurd. For example, it can seem to require an impossible confrontation with a bare universal – the sort of perception imagined in uncharitable readings of philosophic cognition in the *Phaedo* and the *Republic*. Perhaps visual awareness of a property entails seeing a particular. Perhaps visual experience is particular in that sense. But it doesn't follow that visual awareness consists only in seeing particulars, or that visual awareness of properties should be understood on the model of seeing particulars. Rather, the idea is that we're visually aware of the general ways things are, ways which other things could be, in addition to being visually aware of the particular things which are those ways. This idea is compatible with various metaphysical theories about what properties fundamentally are. It is silent on the question whether having a property is instantiating a universal, being a member of a set, comprising a trope which is a member of a set, or something else altogether.

Some philosophers argue that this visual awareness of the general necessarily amounts to belief or knowledge: to see a property in this sense

[11] An alternative would be to say that seeing a property-instance *I* is just seeing in a way which puts you in a position to identify *I*. You can be in a position to identify something without actually doing so. So on this approach, even if identifying *I* involves being aware of something general about *I*, seeing *I* does not entail *actually* being aware of anything general. This is a genuine alternative, but making it both plausible and precise is not straightforward. If to 'identify' *I* is to name or describe *I*, the proposal has seeing depend implausibly on the subject's vocabulary. More plausibly, to 'identify' *I* could be to discriminate *I* from instances of other properties. But there is some pressure to think this condition on seeing a property-instance both too exclusive and too inclusive. Too exclusive, since phenomenal *sorites* cases suggest that small differences in what we see suffice for discrimination only when aggregated (Peacocke 1989). Too inclusive, in light of evidence that the attending required for discrimination alters our inattentive visual experiences, allowing for discriminations which cut finer than those supported by inattentive experience (Ch.6 below; Stazicker 2011).

is already to see *that* something has that property, where seeing that *p* is treated as a species of knowing that *p* (Brewer 2005; Stroud 2009). Although I do not endorse this view, nothing I say here is prejudiced against it. Note, however, that we cannot capture our explanation of Ebdon's behaviour just by saying he does not see that the ball is brown. When he knows which ball is brown, he does see that the ball is brown: he knows that it is brown by keeping track of its location visually. So if what explains Ebdon's mistake is a failure to see that things are brown, it is a failure to see in a specific way that things are brown. More generally, even if it follows necessarily from *Presence* that *S* already sees that something has *P*, *Presence* specifies a special case of *S*'s seeing that something has *P*.[12]

I doubt that there is a simple phrase in ordinary English which we can use, non-stipulatively, to refer uniquely to the relevant aspect of visual experience. For example, to describe what most of us have and Ebdon lacks, we might naturally say that Ebdon does not see which ball is brown. But this does not unambiguously pick out something which Ebdon lacks even when he knows which ball is brown. Like the expression 'he sees that it is brown', the expression 'he sees which one is brown' can be used quite generally to ascribe propositional knowledge which has a basis in visual experience. When Ebdon does know that the ball is brown, there is a perfectly good sense in which he sees which ball is brown. He knows which ball is brown by keeping track of the ball's location visually.

We might say instead that the ball does not look brown to Ebdon. But again this is not unambiguously true, when Ebdon knows which ball is brown. If we show Ebdon Figure 4, the brown ball will look brown to him. Like all the balls other than the red, the brown is in the starting-position dictated by the rules of the game. So the way it looks to Ebdon informs him that it is brown. If there is a use of the expression 'things don't look brown to Ebdon' which captures the relevant feature of his psychology, it is a specific use of that expression.[13]

I think the best thing to say is that Ebdon does not see the specific brown colour of the ball, which most of us would see. I will sometimes use the phrase 'sees *P*' to refer to the aspect of experience on which I want to focus: the aspect of someone's visual experience such that we must

[12] Dretske (1969) attempts to specify the relevant privileged sort of seeing *that*, or 'primary epistemic seeing'. His analysis is given in terms of how things look, so it is beholden to the issue about a specific use of 'looks' which I raise below.

[13] Jackson (1977) attempts to specify a sense of 'looks' which refers uniquely to the relevant special case. See Martin 2010 for criticism.

mention *P* in an adequate description of her visual experience. However, I accept that this usage may involve a technical stipulation, rather than being demanded by the semantics of the phrase as it is used in ordinary English. As it is ordinarily used, talk of seeing colours, shapes and so on may be much more flexible.

Even if there is no phrase of ordinary English which refers uniquely to the relevant aspect of visual experience, can we give it a precise definition using theoretical resources? I am not sure whether we can. The chief difficulty is as follows. *Presence* specifies a distinctive way in which we experience properties, such that we must mention those properties in order faithfully to describe our visual experiences. For example I suggested that, in order to describe the relevant difference between most of us and Ebdon, we should mention the specific brown colour of the ball in describing most people's visual experiences, but withhold mention of it in describing Ebdon's. This proposal trades on the thought that, among our experiences of properties, only some experiences are distinctively perceptual (in this case, visual). After all, when Ebdon knows that the ball before him is brown, he may make a conscious judgement to that effect. To make *Presence* fully explicit, we would need to make precise the distinction between being conscious of a property in this way, and experiencing that property visually.

There are various characteristic features of visual experience to which we might appeal, to try to make the distinction precise. It will be useful to have three of them sketched in outline, albeit very briefly:

(1) It is sometimes said that it's not up to us what we experience visually, whereas it is up to us what to judge or come to believe. For example, Colin McGinn writes:

> The essence of perception is the way in which the world takes hold of one's consciousness, intrudes upon it ... In belief formation there is no analogue for opening one's eyes and having reality flood in.
>
> McGinn 1999: 323

Now visual experience is not entirely passive: seeing involves active exploration of the environment (Noë 2004). Nor is the stream of conscious thought entirely under our control. Perceived events and unsolicited preoccupations play a large part in dictating what you think of, and so in

dictating what conscious judgments you make. Furthermore, you can't just choose on a whim to judge that something is true; in general, you have to see some reason for believing what you judge to be the case.

Nonetheless, McGinn's description reflects a relevant feature of visual experience. What we experience visually is not subject to revision, in light of our other beliefs, in the way in which judgements and beliefs themselves are. The standard example is the Müller-Lyer lines, which (it's claimed) visual experience persists in representing as unequal, even when you know that they are equal. Christopher Peacocke suggests that we can exploit this feature to identify cases of genuinely perceptual experience:

> [T]he characteristic feature of the content of experience as opposed to the content of judgement [is that the former] need not alter when additional information results in a judgement of a content incompatible with the content of the experience.
>
> Peacocke 1986: 156

I think this 'characteristic feature' does provide a useful heuristic, in sorting conscious judgement from visual experience. However, it would be another thing altogether to make the distinction between the two kinds of experience precise in these terms. The feature in question is not sufficient for visual experience. Knowingly or unwittingly, we do sometimes consciously judge contradictory claims to be true. The 'characteristic feature' is probably not necessary for visual experience either. At least, to say that it is necessary is to commit to the controversial empirical claim that visual experience is not cognitively penetrated – roughly, that the processes responsible for our seeing as we do operate independently of what we believe.[14] For if your seeing a property *Q* depends in this way on your belief that *p*, and your belief that *p* gets revised in light of a judgment incompatible with *p*, we should expect that the visual experience will likewise be revised – such that you cease to see *Q*. So it is not at all clear how this 'characteristic feature' of visual experience could be transformed into a precise account of the difference between visual experience of a property on the one hand, and conscious judgement about that property on the other.

[14] For the classic debate, see Churchland 1998 and Fodor 1988.

(2) Intuitively, we might say that the colour of the ball does not figure *immediately* in Ebdon's experience. So we might hope to capture what's distinctive of genuinely visual experience of a property in terms of this notion of immediacy. But this notion is itself hard to make precise.

One tempting thought is that immediacy amounts to an absence of inference. We might say that you experience a property visually if and only if you are thereby in a position to know without inference about that property. However, it is doubtful that Ebdon performs a conscious inference each time he recognises the brown. When things are going well, he may recognise it unreflectively. We might appeal instead to a non-conscious or unconscious inference. But in all of us, the visual system must perform non-conscious computations in order for us to see the colour of the ball. And it is not clear by what criterion we could distinguish the inference we wish to ascribe to Ebdon from computations of this kind.

Figure 5

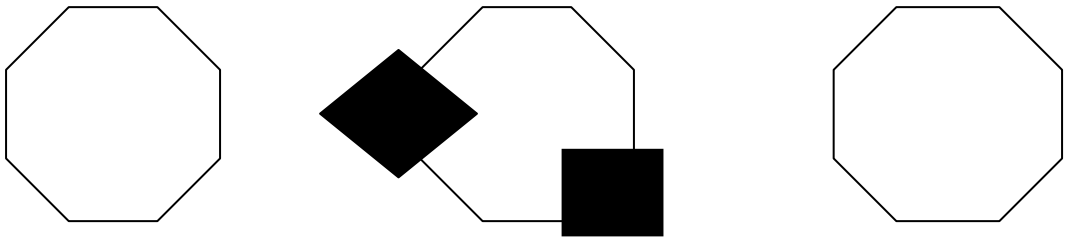

Attempts have been made to distinguish between these different sorts of inference or computation. Gaetano Kanizsa (1985) argued that 'seeing and thinking function according to different rules'; the processes responsible for the 'genuinely perceptual presence' of features of the environment, and by contrast the processes responsible for reasoning, are governed by different principles. For example, consider Figure 5. Kanizsa argued that inductive reasoning would lead us to suppose that the partly occluded shape is similar to those which flank it; by contrast, perceptual processes lead us to experience it as similar to the shape in Figure 6. If Kanizsa's arguments are successful, they show that distinctively 'non-ratiomorphic' principles are among those which govern the processes

responsible for visual experience. But it does not follow that there are no principles which govern both reasoning and visual processing. In fact, that would be hugely surprising. To the extent to which the principles which visual processing exploits are effective ways to discover what the environment contains, effective reasoning is likely sometimes to mirror them – and also vice versa. So here again we have at most a characteristic feature of genuine visual experience, rather than a way of making the distinction between visual experience and conscious judgement precise.

Figure 6

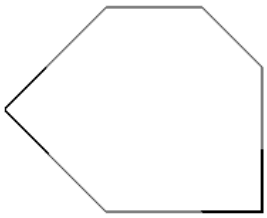

We might try saying that the distinctively inferential basis of Ebdon's knowledge would come out under questioning: if pressed, Ebdon would respond to questioning about how he knows the colour of the ball by citing the path it has followed from its starting-position. But it's not at all clear that counterfactuals about what we would say distinguish the position of most people from Ebdon's position, in a way which generalises into a condition on genuine visual experience of a property. Could there not be someone in Ebdon's position who would fail to cite the path of the ball under questioning? Could the tracking process not become so automatic that it goes wholly unnoticed? And won't many people who do not suffer from Daltonism offer considerations about the lighting, the working order of their eyes, and so, once we question them about how they know? So what, in principle, distinguishes the counterfactuals about what they would say from the counterfactuals about what Ebdon would say?

Perhaps the condition should be normative, rather than merely counterfactual: at least if his belief about the colour of the ball is justified, Ebdon should respond to questioning about how he knows the colour of

the ball by citing the path it has followed from its starting-position. However, to develop this suggestion into a condition on genuine visual experience of a property, we would need to defend some controversial claims in the epistemology of perception. If Ebdon is in general reliable in recognising which ball is brown unreflectively, why should it be a condition on his belief's being justified that he would reconstruct his reasoning in this way? And if there is such a condition on Ebdon's belief's being justified, why does a parallel condition not apply to the rest of us, such that under questioning we must cite the lighting conditions, the working order of our eyes, and so on, if we are to have justified beliefs about the colour of the ball?

(3) A different approach focuses on the logical form of our beliefs about properties. Sometimes we think about the properties we see in a distinctive way, which we can express by talking about '*that* colour', '*that* shape', or '*that* length', and so on. By a Fregean criterion, these thoughts are distinct from our other thoughts about these properties. You might know that something has a length which you would describe as 'that long', without knowing the thing's length in inches, or in any other units of measurement. Similarly, you might know that something has a colour which you would describe as 'that colour', without knowing either a name for the colour or its coordinates in a standard colour space.

This suggests that thoughts which we express in this way involve distinctive modes of presentation of their content, in virtue of which they differ from non-demonstrative thoughts about properties (Peacocke 1989). Furthermore, when you think of a property in this way, you seem to select in thought the very property which you experience visually: to think about *that* colour or *that* shape, you attend to the colour or shape as it figures in your visual experience. So we might take it to be distinctive of visual experience of a property *P* that the experience put one in a position to have thoughts about *P* which have this distinctive form. In that way, we might hope to make precise the idea of distinctively visual experience of a property, as it appears in *Presence*.[15]

However, this proposal is both controversial and hard to make precise. Suppose we're granted the Fregean criterion on thoughts with a

[15] Cf. Snowdon's (1992) account of what it is directly or immediately to perceive an object – although Snowdon says that the direct perception of properties requires a different treatment.

distinctive form: where two thoughts differ in cognitive significance, they differ in mode of presentation. Which thoughts, exactly, have the cognitive significance that's criterial for visual experience of *P*? Not every thought which we express by saying 'that colour' or 'that shape' has the relevant cognitive significance. 'That' can be used anaphorically or in place of a description. Ebdon, for example, might use 'that colour' to refer to whatever colour the ball he sees has, whether or not that colour meets the condition of *Presence* for him.

The obvious suggestion is that your thoughts have the relevant cognitive significance if and only if you can discriminate *P* from other properties. But as we'll see in Chapter 5, this quickly leads to contradiction. For someone with normal human vision, there are three possible colour-patches, *a*, *b* and *c*, related as follows: *a* is indiscriminable from *b*, and *b* is indiscriminable from *c*, yet *a* is discriminable from *c*. The present suggestion has it, incoherently, that you experience *a* as both identical in colour with *c* and distinct in colour from *c*. That cannot be right, so this cannot be the right way to spell out the cognitive significance criterial for visual experience of a property.

To be clear, I don't believe that the cognitive significance of a thought about a property is in fact limited to the property-discriminations which it grounds (see §5.5 below). But I don't see any other obvious way to isolate the cognitive significance which is supposed to be criterial for visual experience of a property. So while it is true that visual experience of a property characteristically puts one in a position to form beliefs we express by talking about 'that colour', 'that shape' and so on, it's not obvious how we could develop this into a precise condition on visual experience of a property.

The above is of course an extremely brief review of the difficulties here. Maybe one of the proposals sketched could be developed into necessary and sufficient conditions on genuine visual experience of a property. And there are also various further possibilities which we could investigate. For example, specific philosophical theories of visual experience give us various proposals about its distinctive nature: it's relational, or it's representational without being propositional, or its representational content is necessarily connected with a distinctive phenomenology, and so on. But as I said, I would like what I say here to be independent of any such specific theory.

In any case, I think we need to take seriously the kind of visual experience which *Presence* specifies, whether or not we can give it a completely precise definition. Progress would be rare in the philosophy of mind, if we had first to make completely precise every distinction to which we appeal. In general, the idea that conscious experience has a certain phenomenal or qualitative character has eluded precise definition. Perhaps we should not be surprised if some specifications of that idea, such as the idea that we enjoy a distinctive kind of visual experience of properties, likewise elude precise definition.

In what follows I will assume that there is an important distinction between visual experience and conscious judgement, and that some of the colours and spatial properties of things meet the condition of *Presence* for us. I will argue that the properties meeting this condition are often only determinable properties. In that sense, visual experience is indeterminate.

Alva Noë's work on 'presence in absence' (2004) has done a great deal to undermine the idea that visual experience presents us with a maximally determinate visual field or visual array. He argues, in part, by appeal to the phenomenon of *amodal occlusion* which I introduced above with reference to Kanisza: the experience of seeing often involves indeterminate information about hidden parts of a scene. Although Noë does not explicitly discuss the experience of determinable properties, his work is obviously a big influence on mine. However, my aims here are different from Noë's. He exploits amodal occlusion and related phenomena to argue for a specific theory about the nature of visual experience: visual experience consists in the exercise of sensory-motor knowledge. My aims are less ambitious but more catholic. I don't argue for a specific theory of the nature of visual experience. Rather, I aim to make a case for visual indeterminacy which should be accepted by adherents of a range of theories about the nature of visual experience.

The phenomena Noë discusses are clearly relevant to *Determinacy* as I formulated it – to the assumption that determinable properties never meet the condition of *Presence* for us, without determinations of those properties also meeting that condition. Consider Figure 7 overleaf. The two horse-shoe shapes are the same size. But the one on the right looks larger, because it seems to have parts which are occluded by the rectangle. Furthermore, in looking to have these occluded parts, the shape looks to have parts with some occluded shape, and some occluded colour. Yet you are not aware of the determinate shape or colour hidden behind the

rectangle. So amodal occlusion might be taken as a counterexample to *Determinacy*.[16]

Figure 7

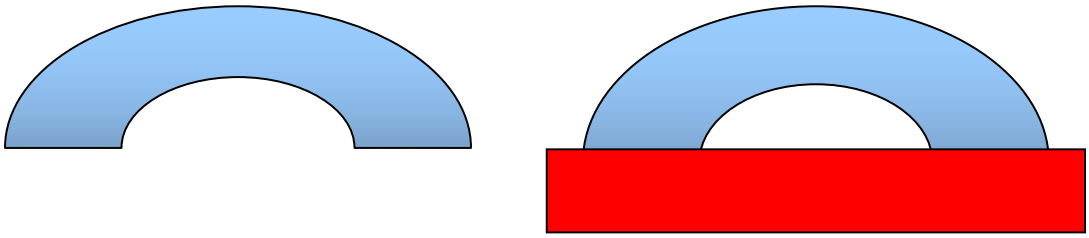

However, for my purposes this is a controversial counterexample, because absent a completely precise definition of *Presence*, it is uncertain whether the occluded shape and colour meet this condition; it is uncertain whether the occluded shape and colour figure essentially in the distinctively visual aspect of your overall experience.

By the criteria sketched under (1) and (2) above, the occluded parts of the horse-shoe are objects of genuine visual experience: their appearance persists even when you know that it is deceptive; the principles by which such appearances are constructed differ from those characteristic of cognition proper (Kanisza 1985). So one might argue (although this move would itself be controversial) that the shapes and colours of those parts are likewise objects of genuine visual experience.

However, it is doubtful that these properties are objects of genuine visual experience by the criterion sketched under (3). First, it is doubtful that you are in a position to form beliefs about the occluded shape or colour which we should express in distinctively demonstrative terms. If you take your experience at face value, perhaps you will believe that the horse-shoe has some occluded shape and colour. But it does not follow that you believe anything about these occluded properties which we should express using the phrases 'that shape' or 'that colour'. Indeed, it's natural to express the commitments of your visual experience with respect to the occluded parts in phrases which are existential, rather than demonstrative, in form: it seems that there exist occluded parts, shapes

[16] Roughly this line of argument was suggested to me by MGF Martin.

and colours belonging to the horse-shoe, not that *this* shape or *that* colour is instantiated there.

Second, suppose we do allow that demonstrative expressions might capture your commitments with respect to the occluded shape and colour. If we do so, we lose the connection between these expressions, on the one hand, and on the other hand the idea that shapes and colours figure in visual experience in a way which enables you to select them in thought. If I ask you to attend to, and so select, the occluded shape or colour, the correct response might be to say that you cannot attend to or select it, because it's hidden from view.

In cases of amodal occlusion, information about hidden parts of a scene does make a difference to how things look. Nonetheless, cases of amodal occlusion are uncertain cases of *Presence*. Someone who takes a restrictive view of the distinctively visual aspect of experience may deny that occluded shapes and colours figure in that aspect of experience. Now Noë argues for an analysis of visual experience according to which occluded parts of a scene figure in visual experience in the same way as unoccluded parts do: they are objects of sensorimotor knowledge. If Noë's theory is correct, then the restrictive view of distinctively visual experience is a mistake. However, I aim to demonstrate that *Determinacy* is false independently of any specific theory of visual experience, and independently of how liberal or restrictive a view we adopt of the contents of visual experience. So I will set aside cases of amodal occlusion, and focus instead on cases in which the condition of *Presence* is unequivocally met, assuming that it's ever met. I think the argument below should be accepted even by someone who adopts a maximally restrictive view of the contents of visual experience.

2.2 *Indeterminacy*

It will be convenient to read *Determinacy* in representationalist terms. So understood, *Determinacy* states that visual experience never represents determinable properties, without representing determinations of those properties. In this part of the chapter, I will argue that *Determinacy* is false with respect to the colours and spatial properties of things. That is, I will argue for:

Indeterminacy

> Visual experience often represents determinable colours and spatial properties, without representing determinations of them.

With respect to spatial properties, I will also argue for the stronger, because universal, thesis that the spatial properties which visual experience represents are *always* determinable.

First we need to understand what it is for a property to be determinable. Determinability is not the same as inexactness. Rather, determinability is a species of inexactness. For example, if I tell you that my shirt is either blue or long, I have told you something inexact about the shirt. But being either blue or long is not a determinable property.

Similarly, it is often said that to have a property which determines a property *P* is to have *P* in a specific way. But to be red and square, say, is in a perfectly good sense to be red in a specific way. Being red and square does not determine being red. So this is not a sufficient condition on determination.

In recent work, Stephen Yablo (1992) has used the following purely modal condition on determination. Where '*P*' and '*Q*' express a determinable and its determination respectively:

$$\Box((\forall x)(Qx \rightarrow Px)) \;\&\; \Diamond((\exists x)(Px \;\&\; \sim Qx))$$

Necessarily, everything crimson is red, but possibly – indeed actually – there are red things which are not crimson. However (as Yablo himself notes) this condition is not sufficient for determination as it has more traditionally been understood (cf. Johnson 1921; Prior 1949). For example, the purely modal condition does not distinguish determinations from conjunctive properties. Necessarily, whatever has a conjunctive property has each of the properties conjoined, and it is possible for something to have one of the properties conjoined without having the conjunctive property. Perhaps there are no conjunctive properties. But if so, we need to understand what distinguishes determinations from these spurious entities – or alternatively what distinguishes determinables from (spurious) disjunctive properties.

I will work with a summary version of Eric Funkhouser's (2006) analysis of determination. We can think of a determinable property as having *determination dimensions.* A determination dimension is a variable

along which a property is determined. For example, being coloured has the determination dimensions hue, saturation and brightness. Being coloured is a structurally simple case: to have any determination of being coloured is just to have values within some range for hue, saturation and brightness. Other properties are more complex, in that they're determined along a varying number of determination dimensions. Being (2-D) shaped has determination dimensions for number of sides, length of each side, and size of each interior angle; the number of dimensions for side-length and angle-size is a function of the number of sides. Still, being coloured and being shaped each fix at least the *schematic* determination dimensions I have described (Funkhouser 2006: 556).

Property *P* determines property *Q* if and only if *P* differs in nature from *Q* only along the schematic determination dimensions of *Q*, such that the values along these variables consistent with instantiating *P* are a proper subset of the values consistent with instantiating *Q*.[17] For example, being scarlet differs in nature from being red only in that the range of hue, saturation and brightness consistent with being scarlet is a proper subset of the range consistent with being red. Being triangular differs in nature from being shaped only in that the number of sides and range of side-lengths and angle-sizes consistent with being triangular is a proper subset of those consistent with being shaped. By contrast, being red and square differs from being red in squareness, which is not a value for any schematic determination dimension of being red.

On this analysis, a property *P* may determine many determinables – all those determinables instantiation of which is consistent with a range of determination-dimension values that's a super-set of the values consistent with instantiating *P*. So – importantly for my purposes – where an object has *P*, two veridical representations of the object may represent it as having two different determinable properties.

Indeterminacy says that visual experience often represents only properties corresponding to ranges, rather than absolute values, along the determination dimensions of colours and spatial properties. So for example your visual experience might represent a line as *between 9.995 cm and 10.005 cm long*, rather than as (exactly) *10 cm long*. By this I mean that

[17] This is only the gist of Funkhouser's detailed analysis and mathematical model. As stated here the analysis is obviously circular, albeit informative. It's not clear to me whether Funkhouser's analysis avoids this circularity. That depends on the precise role of his mathematical model.

the length *between 9.995 cm and 10.005 cm* meets the condition of *Presence* for you. However, it does not follow that the upper and lower bounds of the range in question meet this condition – that we must mention the lengths *9.995 cm* and *10.005 cm* in a faithful description of your experience, or that the experience puts you in a position to know about those upper and lower bounds. Neither the units of measurement, nor the limits of the range expressed in any other way, need meet the conditions of *Presence*. Rather, you are aware of a property which we, as interpreters, may describe in terms of its upper and lower bounds. To mark this, we might say that your visual experience represents a property *P*, and *P* is in fact the length *between 9.995 and 10.005 cm*. But to mark it in that way each time would be tiresome.

The clearest empirical justification for *Indeterminacy* lies in the limited and varying *resolution* of vision. Let's focus first on spatial resolution and indeterminacy in the visual experience of space. The spatial resolution of a representation is given by the maximum spatial frequency to which it is sensitive. For standard experimental stimuli, spatial frequency is measured in sinusoidal cycles of variation per degree of visual angle (cpd). In Figure 8, spatial frequency increases from left to right, in that the rate of sinusoidal variation from light to dark and back increases from left to right. Intuitively, we can think of the spatial frequency of a stimulus as its rate of variation across space. For example, where a stimulus has a higher spatial frequency, it typically has a more detailed pattern.

Figure 8

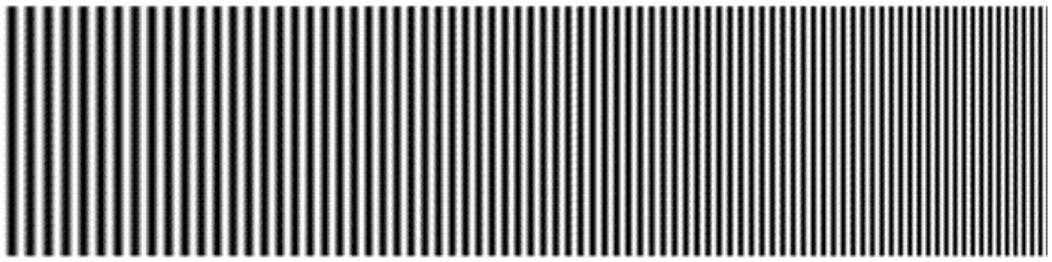

Visual processes respond to distant objects in lower spatial resolution, relative to nearby objects: more rapid variations across space

are visible the nearer the object is.[18] For example, if you walk a little distance away, you won't see the variations on the right of Figure 8. Spatial resolution also decreases away from the fovea at the centre of the eye. Early in post-retinal visual processing, information about the scene you see passes through varying spatial filters, attuned to various spatial frequencies. As the locations of these filters get further from the fovea, they're attuned to progressively lower spatial frequencies: higher-frequency detail goes unprocessed. So peripheral vision has a progressively lower spatial resolution. Spatial resolution is limited even for foveated objects: even when you look right at something, there are spatial frequencies across it to which early visual processing is insensitive (De Valois and De Valois 1988).

Because visual processes are insensitive in these ways to high rates of variation across space, they're insensitive to absolute values along the determination dimensions of spatial properties. Consider an (artificially simple) example. You're presented with a line that's clearly visible, but a fair distance away. Call your visual response to the line's length '*R*'. Conditional on *R*, there is a high probability that the line has any of a range of lengths, say between about 9.995 cm and 10.005 cm. Conditional on *R*, the probability that the line has a length around the edge of this range is lower, and the probability that it has a length much outside the range is negligible. This is a corollary of the fact that, at the distance at which the line is presented, visual responses are not sensitive to variations which occur across less than about 0.01 cm. (Typically, the probability distribution will be roughly Gaussian or bell-shaped. At any rate it is continuous, so there is no sharp cut-off between probable lengths and improbable lengths.)

As a result, if the line is in fact exactly 10 cm long, *R* is not reliably correlated with this exact length. Rather, *R* is reliably correlated with a range of lengths between about 9.995 cm and 10.005 cm. This is consistent with the relationship between stimulus lengths and visual responses being roughly linear. Overall, as the length goes up, the level of visual response does likewise. But for a specific level of response, any length within a range is probable. Note also that the ranges of properties to which different responses are sensitive in this way may overlap. For example,

[18] Strictly speaking, this is true only where spatial frequency is given in terms of a standard measure of distance, rather than degrees of visual angle.

conditional on a response slightly greater than *R*, any of a range of lengths between 9.996 and 10.006 cm might be probable (Figure 9).

Figure 9

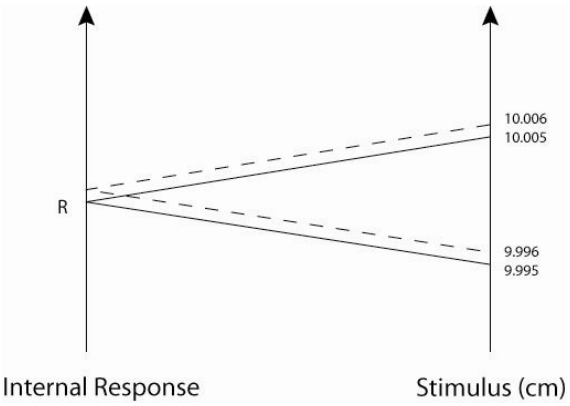

Now it is a basic principle of interpretation that, when we assign semantic content to a representational system, we should do so in such a way that the contents come out by and large veridical (Dennett 1978). This is consistent with the possibility that many representations in the system are not veridical. For example, interpreting vision in this way is consistent with the possibility that attention causes widespread visual illusions. What's required is just a background of veridicality.

This principle is an important heuristic device, irrespective of the approach we take to *constitutive* conditions on the content of visual experience. For example, we should employ the principle even if the content of visual experience supervenes on local states of the brain.[19] Further considerations might ultimately trump this principle, and I will consider some relevant further considerations in §2.3. But absent specific reasons for taking a different approach, we should say that a visual representation *R* represents a property *P* only if the occurrence of *R* is

[19] If visual experience is fundamentally relational, rather than representational, then the condition on visual experience of a property is stronger: in no case you can you stand in the visual relation to a property which is not instantiated before you. Still, we may apply the criterion of by-and-large veridicality to the information-processing states which make it possible for one to stand in the visual relation to a property. Only if these states are by-and-large veridical will they by-and-large perform their function and make relational visual experience possible.

reliably correlated with the instantiation of *P*. The reliable correlation should hold across the range of circumstances such that *R* occurs in a normal way. If *R* is correlated with *P* only under exceptionally ideal or peculiar circumstances, then *R* will not be by and large a veridical representation of *P*. Equally, if *R* sometimes occurs in abnormal circumstances where *P* is not instantiated, this shouldn't count against the claim that *R* represents *P*.

So the notion of sensitivity that's relevant for our purposes is:

Sensitivity

> In every background condition such that *R* occurs in the normal way, if *P* had not been instantiated, *R* would not have occurred.[20]

On a standard semantics for the counterfactual (Lewis 1973), and where the instantiation of *P* is a contingent matter, *Sensitivity* is true iff:

> For every possible world *w* at which background conditions are such that *R* occurs in the normal way, there is at least one world *x* at which *P* is not instantiated and *R* does not occur, such that *x* is closer to *w* than every world at which *P* is not instantiated and *R* does occur.

Which spatial properties and visual representations meet this condition depends on the spatial resolution of vision. For example, take an object with maximally determinate length *L*, seen by someone whose (normally occurring) visual representation of the object's length is *R*. Whether or not the object has *L* does not make a difference to whether *R* occurs in the observer. There are possible worlds arbitrarily close to actuality, such that the object has a length which differs from *L* by an arbitrarily small magnitude. At at least one of these close worlds, *R* still occurs.[21]

---

[20] Note that taking *Sensitivity* to be a condition on visual representation of a property does not commit us to thinking that there is an analogous condition on propositional knowledge, such that knowledge that *p* must be sensitive to the proposition that *p*. See Williamson 2002, Ch.7, for criticism of that idea.

[21] At at least one, because *L* might in principle be at the upper or lower bound of lengths beyond which *R* would not have occurred, in background conditions such that *R* occurs in the usual way.

This gives us good reason to deny that the visual representation *R* represents *L*. On that assignment of representational content, *R* would be veridical only by astonishing accident. Since *L* does not make a difference to whether *R* occurs, there is no reason to expect *R* to be reliably correlated with *L*. By the same reasoning, there are various less determinate lengths which *R* does not represent.

By contrast, there is a determinable length to which *R* is sensitive. In the example above, this might be the length: *between 9.995 and 10.005cm*. There are worlds at which *R* does not occur and the line is longer than 10.005cm by an arbitrarily small magnitude. These worlds are closer to actuality than every world at which *R* does occur and the line is not between 9.995 and 10.005cm long. This gives us a reason to say that, if *R* represents a length, it's the length *between 9.995 and 10.005cm*: *R* is reliably correlated with this length. Since the spatial resolution of vision is always limited, the argument generalises: for all visual experiences, and all spatial properties, visual experience represents only determinable spatial properties.[22] Call this thesis:

Spatial Indeterminacy

> The spatial properties which visual experience represents are always determinable.

With respect to the colours, I will argue for a weaker thesis:

Colour Indeterminacy

> Visual experience often represents determinable colours, without representing determinations of them.

The thesis about colour experience directly analogous to *Spatial Indeterminacy* would be the universally quantified:

[22] We could make a different argument that vision represents determinable properties, if we accepted a causal condition on representation, together with Yablo's (1992) argument that determinable properties are causes, in virtue of the fact that they make a difference to effects. Nothing here turns on that argument.

*Colour Indeterminacy**

> The colours which visual experience represents are always determinable.

For all I know, *Colour Indeterminacy** might be true, but I know of no conclusive argument for it. This is not because visual sensitivity to spectral properties of the light is somehow unlimited or infinite. Limited resolution is a quite general feature of visual processing. Rather, it is because insensitivity to spectral properties of the light need not entail insensitivity to the colours of things.

We can get a handle on the issues here without going into the details of any very specific metaphysical theory about the colours. The issues are clearest if we assume that the colours are response-dependent, in the sense that the identities of the colours depend on human visual responses. So let's start there, before looking at some other broad ways of thinking about the nature of the colours.

One very simple version of this idea about response-dependence is:

Disposition

> For a surface to have a certain colour is for that surface to be disposed to cause a certain experiential response in normal human observers under normal viewing conditions.[23]

On this view, for a surface to have a maximally determinate colour *C* (i.e. to have absolute values for each of hue, saturation and brightness) is for that surface to be disposed to produce some experiential response *E* in a normal human observer, under normal viewing conditions. This leaves open the possibility that *E* might be a determinable experience-type, rather than a maximally determinate one: to have *C* might be to be disposed to cause any experiential response within a certain range. But *Dispositionalism* closes off the possibility that there might be a colour to which humans are insensitive (in the sense specified by *Sensitivity*). For there to be a colour *C* is, *inter alia*, for there to be a type of experience *E* such that, in background conditions such that *E* occurs in the normal way, *E* would not have occurred if *C* had not been instantiated. So at least on

[23] For views in this broad spirit, see e.g. McGinn 1983 and Smith 1990.

this simple version of the view that the colours are response-dependent, *Colour Indeterminacy** is necessarily false – despite the fact that visual responses are insensitive to variations in absolute values of the spectral properties of light.

Things are a little more complicated given different views of the nature of the colours. Consider the view that to have a certain colour is not to be disposed to cause a certain experience, but to be disposed to reflect the light in a certain way:

Reflectance

> For a surface to have a certain colour is for that surface to be disposed to reflect the light in a certain way.[24]

Which dispositions of this kind are colours? According to one version of *Reflectance*, only a disposition to reflect the light in a way to which human visual experiences are sensitive is a colour. A colour *C* is a disposition to reflect the light in a way *W*, such that there is a human experience-type *E* which would not have occurred, if the light had not been reflected in way *W*. On this version of *Reflectance*, there is an experience-type sensitive to every colour, and *Colour Indeterminacy** is necessarily false.

There are also versions of *Reflectance* which do not cast the colours as response-dependent in this way. For example, we might take every way of reflecting the light to fix a colour: a colour *C* is just a disposition to reflect the light in some way *W*. On this version of *Reflectance*, there are colours to which no human visual experience is sensitive, since there are minute variations in how the light is reflected to which no human visual response is sensitive. However, it is still not obvious that *Colour Indeterminacy** follows, because this version of *Reflectance* makes it unclear how determination for the colours is to be understood.

Recall that a colour *C* determines a colour *D* if and only if *C* differs in nature from *D* only along the determination dimensions of hue, saturation and brightness, such that the values along these dimensions consistent with instantiating *C* are a proper subset of the values consistent with instantiating *D*. Why should we suppose that every disposition to reflect the light in some exact way has a proprietary location in this property space? Could there not be two different exact ways *W* and *Y* of

[24] For views in this broad spirit, see e.g. Tye 2000; Byrne and Hilbert 2003.

reflecting the light, such that it is altogether arbitrary to assign one position in colour space to *W* and another position to *Y* – the difference between *W* and *Y* having no principled, relevant connection with any of the differences in virtue of which we experience different values for hue, saturation, or brightness? If so, then the disposition to reflect light in way *W* and the disposition to reflect light in way *Y* are not different determinations of a single determinable colour (as I defined the determination relation for colours). And if some exact colours are not determinations of the determinable *coloured*, it's not obvious how we should understand claims about determinable colours, including *Colour Indeterminacy**.

In fact, there is some reason to believe that there are infinitely many ways of reflecting the light differences between which have no principled, relevant connection with transformations in colour space. Light distributions with various different constituent wavelengths form *metamers* – that is, these light distributions are indistinguishable in colour to a human observer. In particular, there are infinitely many different triples of spectral lights which produce one metamer when the three lights are combined (Hardin 1993: 28). (Triples, because of the trichromatic structure of human vision. The constraint is just that one of the three lights must be mostly blue, one mostly green, and one mostly red.) On the face of it, there could be no principled reason to represent the difference between any two such triples by any particular transformation along the dimensions of hue, saturation and brightness. The determination dimensions of colour just have the wrong structure to capture this difference in the exact properties of the light.

At this point you could respond in a couple of different ways, if you wanted to defend both *Reflectance* and *Colour Indeterminacy**. You could try to preserve the organization of the colours in their traditional determination space. For example, you could propose the following specific version of *Reflectance*: to have a colour *C* is to be disposed to reflect the light in a way *W*, where to reflect the light in way *W* is to reflect it in any of a set of ways the differences between which have no principled connection with differences in hue, saturation or brightness. According to this proposal, many different ways of reflecting the light fix just one colour, but the colours cut finer than metamers. The colours cut finer than metamers, because there are some differences in the light which we would experience as differences in hue, saturation or brightness, if only our eyes were more sensitive along those dimensions. These differences in the light

do have a principled connection with differences in hue, saturation or brightness. So every colour does determine the determinable *coloured*. And on this view there are determinate colours to which human visual experience is altogether insensitive.

Alternatively, you could propose a revision of our naïve view about determination dimensions of colour. You could argue that the real determination dimensions of colour are not hue, saturation and brightness, or anything of their ilk. Rather, the determination dimensions of colour are to be given in terms of response-independent, physical properties of the light. Presumably there will then be determinate colours to which human visual experience is altogether insensitive.[25]

I don't know whether either of these responses is plausible, and I will offer no argument for or against them. Suffice it to say that these are very specific versions of *Reflectance*. Even assuming *Reflectance*, *Colour Indeterminacy** does not follow in any immediate way from the fact that human visual experience is insensitive to minute variations in the spectral properties of light.

Finally, consider a third broad view about the nature of the colours:

Primitivism

> The colours are intrinsic properties of surfaces. Only topic-neutral and chromatic vocabulary figures essentially in statements of their nature.[26]

For example, 'Every shade of orange is reddish' expresses the nature of colour; 'Surfaces reflect light' and 'Surfaces cause colour-experiences' do not. *Primitivism* posits no constitutive connection between the colours and either the light or visual experiences. On this view colours, like shapes, are

[25] This position is in the spirit of Kalderon 2007. He argues that every experienced colour is a determinable of some physical determinate. Although Kalderon does not explicitly address concerns about the determination dimensions of colour, he does claim that the structure of traditional colour-space is response-dependent, while the colours themselves are not. Yablo (1995) likewise argues that colours are 'nonphysical determinables' of 'microphysical determinates'. Yablo does not worry about the determination dimensions of colour, but this is because his notion of 'determination' is really just the more general notion of inexactness (see p.47 above).

[26] For this way of specifying primitivism, see Byrne and Hilbert 2003. For a view in this broad spirit, see Campbell 1993.

just intrinsic properties of the visible environment. So *Primitivism* gives us no reason to believe that *Colour Indeterminacy** is necessarily false: it gives us no reason to believe that there could not be colours to which human visual experience is insensitive. But equally, if *Primitivism* is true, *Colour Indeterminacy** does not follow from the fact that human vision is insensitive to minute variations in the light. *That* insensitivity is consistent with there being only colours to which human vision is sensitive, because differences in colour are all gross enough to make a difference to human visual responses.

There are, of course, many further possible views about the nature of the colours. But the above shows, I hope, that the general limits on human visual sensitivity do not suffice to establish *Colour Indeterminacy**.[27] So let's agree, for the sake of argument, that there is some human experience which is sensitive to every colour.

Even so, we should accept the weaker *Colour Indeterminacy*:

> Visual experience often represents determinable colours, without representing determinations of them.

Evidence for *Colour Indeterminacy* lies in the fact that the resolution of human colour vision is *variable*. Even if, under ideal conditions, the visual system is sensitive to maximally determinate differences in colour,

[27] Hellie (2005) argues that human visual experience represents only inexact colours, even under ideal viewing conditions. As far as I can see, the problems just described are serious problems for his argument. The inexactness Hellie canvasses is inexactness only along the dimensions of hue, saturation and brightness. So in my terms he argues for the more specific thesis that human visual experience represents only determinable colours. Hellie argues that, given noise in visual processing, there is no reliable causal connection between a maximally determinate colour and a type of colour-experience; therefore, assuming a causal-nomological theory of representation, no experience-type has a maximally determinate colour as its semantic value. (Noise is one source of limited resolution in human vision, the absolute insensitivity of receptors another. It's not clear to me why Hellie focuses only on noise.) Now noise in visual processing does prevent reliable connections between experience-types and maximally determinate wavelengths at the retina. But as we've seen it does not follow that this noise is sufficient to prevent reliable connections between experience-types and absolute values for hue, saturation and brightness. Hellie also argues that 'phenomenal noise', the 'flicker' which makes colour-experience unstable, occludes maximally determinate colours. It's essential to his account that this noise consist in a sort of sense-datum, rather than an illusory representation of fluctuating maximally determinate colour. I'm not sure whether there are such sense-data, or whether it makes sense to suppose that they occlude colours.

conditions are often not ideal, and the sensitivity of colour vision varies. If *Sensitivity* is a condition on visual representation of a property, the determinacy of the colours which visual experience represents likewise varies.

First, it is well known that we are less sensitive to colours in peripheral vision than to colours in foveal vision (Moreland and Cruz 1959). Systematically, subjects fail to discriminate hues at peripheral locations, when they can reliably discriminate those same hues at fixation (Noorlander et al. 1983). Like the limits on peripheral spatial resolution, these limits on peripheral colour resolution have their source deep in the visual system. According to some authors, peripheral cells in the retina are less sensitive to variations in hue than foveal cells are (Shapley and Perry 1986). More recently it's been argued that peripheral cells in the retina are just as sensitive to hue as foveal cells; the source of insensitivity lies in the post-retinal visual cortex instead (Martin et al. 2001; Solomon et al. 2005). At any rate, the limits on your sensitivity to colours at peripheral locations are not simply limits on what you can judge, given visual experience which is constant across locations. They are limits on the sensitivity of visual experience. At foveal locations, visual experience is sensitive to colours more determinate than those to which it is sensitive at peripheral locations. If *Sensitivity* is a condition on visual representation of a property, at peripheral locations visual experience represents determinable colours without representing determinations of them.

There are also further features of visual processing in virtue of which sensitivity to both colours and spatial properties is limited, under common sub-ideal conditions. One striking example is the *crowding effect*:[28] when a stimulus which subjects can usually identify is presented slightly away from fixation, surrounded by other slightly different stimuli, the subject is sometimes unable to identify the stimulus. Herman Bouma (1970) first demonstrated this effect on the experience of alphabetic characters. He found that subjects can see but not identify letters under the conditions described, even though they could readily identify those same letters under other conditions. More recently, a statistically similar effect has been demonstrated for the orientation, size, saturation and hue

[28] This phrase is often used to refer to a rough family of phenomena with somewhat various statistical properties. Strictly speaking, the crowding effect should be distinguished from *lateral masking* (Levi 2008). But for our purposes the rough usage will suffice.

of a stimulus: under the conditions described, subjects cannot tell whether a stimulus matches a sample for orientation, size, saturation or hue; under other conditions they can tell this (van den Berg et al. 2007). If you look at Figure 10, you might experience the crowding effect for yourself. For each panel, fixate the central cross and try to say whether the central item on the left matches that on the right.

Figure 10

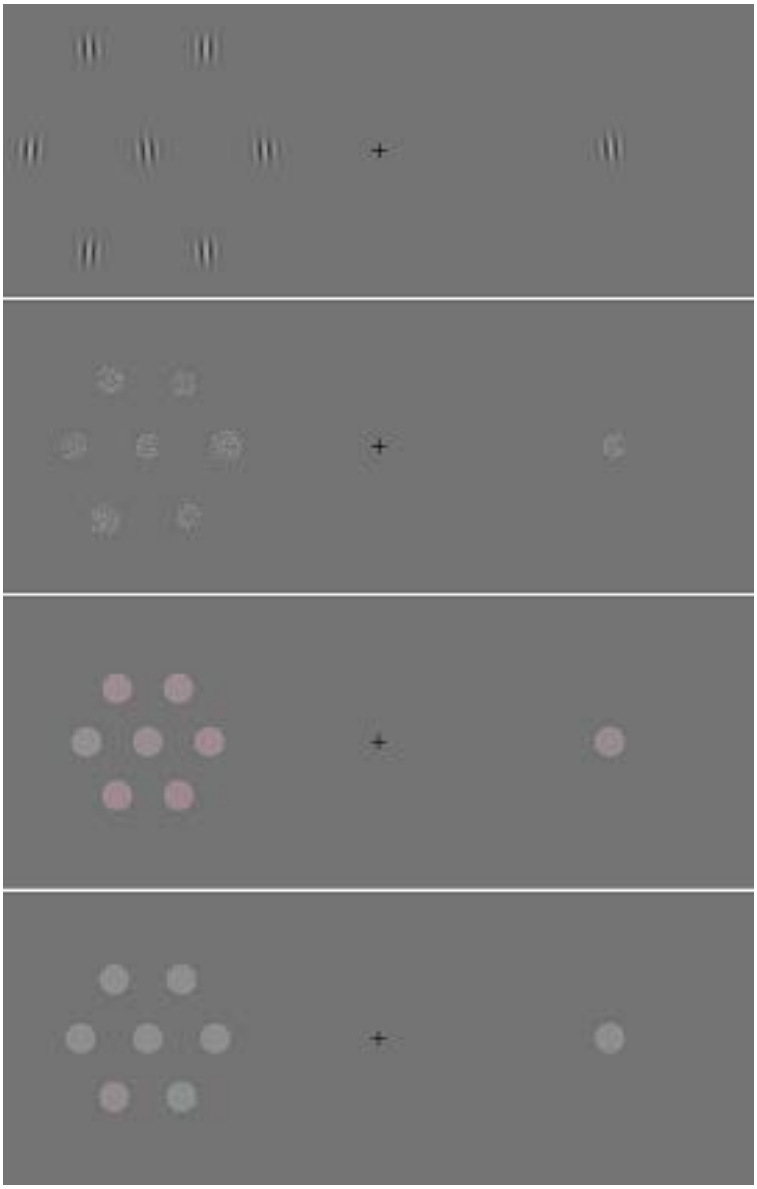

Reproduced from van den Berg et al. 2007: 2, by permission of the Association for Research in Vision and Ophthalmology (copyright 2007).

Just introspectively, the crowding effect seems to be an effect on visual experience: you can see the colours of the stimuli on the left, but you cannot see them determinately enough to see which of them best matches the stimulus on the right. This is confirmed by the fact that the mechanisms of crowding lie deep in the visual cortex, perhaps even as early as V1 (Levi 2008). In conditions which do not produce the crowding effect, visual experience is sensitive to properties more determinate than the properties to which it is sensitive under the conditions of crowding. The conditions of crowding often obtain in a natural environment: similar items are often grouped together, and you don't always look directly at them. So if *Sensitivity* is a condition on visual experience of a property, visual experience often represents determinable properties without representing determinations of them – and these properties include the colours, even if under ideal conditions visual experiences are sensitive to maximally determinate colours. That is, if *Sensitivity* is a condition on visual experience of a property, then both *Indeterminacy* and *Colour Indeterminacy* are true.

Now I said that *Sensitivity* should fix our interpretation of a representational system, absent specific reasons for taking a different approach. For some representational systems there might be such reasons. Take the case of natural language. In some contexts though not in others, you can truly assert that France is hexagonal (Austin 1962). On one (controversial) way of understanding this, the semantic content asserted remains the same in these different contexts; what changes is the degree of precision required for truth. On this approach, the assertion that France is hexagonal represents a property to which it's not sensitive: the property of being strictly-speaking hexagonal. The assertion may nonetheless be veridical, if veridicality need not be *precise* veridicality. You might motivate this approach by noting that the term 'hexagonal' is shared between this assertion and others – for example the assertion that a shape is hexagonal if and only if it's a closed plane figure with six sides. If the latter assertion fixes an interpretation for the term 'hexagonal', and we assume that semantic content supervenes on linguistic expression, then perhaps we have a principled reason for taking the assertion that France is hexagonal to represent a property to which it is not sensitive.

There is no parallel reason for taking visual experience to represent colours or spatial properties to which it's not sensitive – and so for denying *Indeterminacy*. There is no identifiable vehicle of visual representation that's shared between different cases in a parallel way.

However, you might propose a different reason for taking visual experience to represent spatial properties to which it's not sensitive. Recall that, where I say that visual experience represents a property *P*, I mean that *P* figures in the phenomenology of visual experience: a faithful description of your visual experience must mention *P.* And it is tempting to think that this constraint gives us reason to suppose that visual experience represents properties to which it's not sensitive. It is tempting to think that visual phenomenology consists in being presented with maximally determinate spatial properties. Next, I will argue that this temptation is a mistake.[29]

2.3 *Indeterminacy* and Phenomenology

The properties which your visual experiences represent are determinable, and these properties figure in the phenomenology of your experiences: we need to mention these properties if we are faithfully to describe the character of your visual experiences. Why should we think this implausible, as compared with the claim that maximally determinate properties figure in the phenomenology of visual experience? For example, to represent the maximally determinate spatial properties of an object is to specify those properties in terms of spatial points.[30] Why would one think, upon phenomenological reflection, that visual experience represents these properties, or that these properties figure in the phenomenology of visual experience? You might just think that's introspectively obvious. But it is a general, theoretical claim, and we should be suspicious of appeals to introspection to justify such claims in an immediate way. To motivate the claim, we would need careful phenomenological descriptions of particular cases. And we can't expect people to describe the shapes they see in terms of spatial points. In fact, we usually think descriptions of visual experience under-specify its content.

---

[29] In general, another reason for taking representational content to outstrip *Sensitivity* lies in the view that a representation's content is fixed by its evolutionary function (Millikan 1984). On this approach, a representation may be systematically falsidical, where the evolutionary advantage of its being so outweighs the cost. I take it there would be no such advantage in representing determinate colours or spatial properties falsidically.

[30] A spatial point is in principle indivisible. In some metaphysical accounts space is gunky: there are no spatial points. Nothing here rests on that view.

I suspect that the lingering role of *impressions* in our thinking is an important source of the idea that visual phenomenology presents us with maximally determinate properties. As I said, if we accepted Hume's theory of impressions, we would have a theoretical reason to say that where a determinable property is present to you, a determination of that property is also present to you. Few today accept the theory of impressions, but talk of impressions persists.[31] And once we reify visual experience in this way, it's tempting to confuse determinacy in experience with the determinacy of its objects. It's tempting to think that, if the things we see have maximally determinate spatial properties, visual experience must represent such properties. In principle, you could reify visual experience without drawing this conclusion, since determinate properties of an image may represent less determinate properties of reality. (Think of shading in a sketch.) However, in practice talk of impressions tends to obscure the possibility that visual experience represents determinable properties.

If we lack *positive* phenomenological evidence that visual experience represents maximally determinate spatial properties, perhaps a more promising strategy is to argue negatively, to object that *Indeterminacy* somehow gets the phenomenology wrong. One such line of objection is as follows: 'Representing a determinable property is equivalent to representing the disjunction of that property's determinations. It's phenomenologically implausible that visual experience represents disjunctions, so it's implausible that visual experience represents determinable properties.'

Now if representational content is given extensionally, in terms of the possibilities logically compatible with it, it's true that representing determinables is equivalent to representing disjunctions of their determinations. But this approach to representational content gives up the aim of capturing phenomenology in the structure of the contents we ascribe. This shows up in the fact that, on such an approach, representing any content *C* is equivalent to representing the disjunction of *C* and a necessary falsehood. So this line of argument cannot be used to motivate

[31] Some examples which will prove notable in later chapters are Block 2007 and Carrasco 2004. I should note that for Hume himself, the relevant maximal determinations are *minima sensibilia*, the most determinate properties discernible by our senses (*Treatise* I.1.ii). So Hume could allow that spatial points are not visibly present, where point-hood is measured according to a physical standard.

the view that visual experience represents maximally determinate properties. Such representations are disjunctive by the same criterion.

Furthermore, on the analysis I'm working with, there is a distinctive metaphysical relation connecting determinables with their determinates, which does not connect between mere disjunctions with their disjuncts. As a result, where several possibilities all involve instantiation of the same determinable property, these possibilities are unified by more than mere disjunction. So even if we use these possibilities to capture visual contents, there is a unity in representations of determinables which is lacking in merely disjunctive representations.

We can understand this unity as having a phenomenological dimension, in the following way. Suppose we're using phenomenology as a guide to the representational contents we assign to someone's visual experience. What must her experience be like, subjectively, for us to say that its content is a set of possibilities each of which involves instantiation of the same determinable shape? Well one idea is that her experience has an explicitly disjunctive import: just through the experience, she understands that the object she sees has one of a set of different determinate shapes, and she understands what those possible determinate shapes are. But there is an alternative to this unrealistic idea. Given that determinables are not merely disjunctive properties, we need not find any explicitly disjunctive import in the subjective character of experience, in order to assign the content in question. We can assign this content on the basis that experience seems to make apparent a single, unified shape, without making apparent which maximally determinate shape it is. Experience may then be silent, in its subjective import, on the more specific disjoint possibilities consistent with the object's having this single, unified shape.

Compare a different sort of case. If Liz thinks that Pierre is in France, we can model the content of Liz's thought as a set of several possible worlds – worlds at which Pierre is in Paris, worlds at which he's in Lyon, worlds at which he's in the Camargue, and so on. In doing so, we don't require that these more specific possibilities be subjectively salient to Liz. She might not know of the Camargue at all. What justifies the model is that France is subjectively salient to Liz, and that France in fact includes these locations. Similarly, visual representation of determinable properties need not introduce disjoint possibilities into visual phenomenology. Visual experience of a determinable property may have the phenomenology characteristic of experience of a single, unified shape.

It is also important to distinguish representation of determinable properties from some other forms of representational indeterminacy. For example, John Searle (1987) finds it phenomenologically implausible that the contents of conscious episodes should be indeterminate. But what's objected to here is the idea that conscious contents might be subject to *Quinean* indeterminacy about reference. This is quite different from the indeterminacy for which I have argued. Quine's thought was, very roughly, that there is no principled reason for saying that a cognitive response represents one aspect of what stimulates it rather than another; the response does not determinately represent any aspect of the stimulus. By contrast, the indeterminacy to which I'm appealing falls within the scope of the representation: a state that is determinately a representation of blue represents a determinable property; in that sense, the representation is indeterminate between royal blue and navy blue. My proposal is consistent with the idea that there's a determinate fact about the content of every conscious state.

One might similarly object to the idea that conscious visual representations are *vague*. A representation is vague where its content is not fixed precisely, so that there may be no determinate fact about whether it represents things as they are. For example, you might claim that it is indeterminate whether visual experience represents a line as between 9.995 cm and 10.005 cm long, or as between 9.997 cm long and 10.003 cm long. If the line is in fact 9.996 cm long, it will then be indeterminate whether visual experience represents things as they are. Call this proposal:

Vague Indeterminacy

> It is indeterminate which of several determinable properties a given visual experience represents.

Now my approach is compatible with, but does not require, this sort of vagueness in visual representation. If *Vague Indeterminacy* is phenomenologically implausible, that does not constitute an objection to what I have said, since neither *Indeterminacy* itself nor my argument for it requires *Vague Indeterminacy.*

I accept that there may well be vagueness in which determinable properties a visual experience is sensitive to. As I said, the probability distribution linking a visual response with visible properties involves no

sharp cut-off between probable and improbable properties of a stimulus, conditional on the visual response. There may be no precise probability beyond which that we should say that the visual response would not have occurred, had the property not been instantiated. As a result, the upper and lower bounds of the most determinate property to which a visual response is sensitive may not be fixed precisely. For example, it might be indeterminate whether the most determinate length to which a response is sensitive is *between 9.995 and 10.005 cm* or *between 9.997 and 10.003 cm.*

Nonetheless, *Vague Indeterminacy* does not follow from taking *Sensitivity* to be a condition on visual representation of a property. First, recall that *Sensitivity* concerns only what would happen in circumstances such that a visual response *R* occurs in the normal way. Here it is again:

Sensitivity

> In every background condition such that *R* occurs in the normal way, if *P* had not been instantiated, *R* would not have occurred.

This condition of normalcy excludes cases in which there is a very low probability that *P* is instantiated, conditional on *R* (and *vice versa*): *R* may be sensitive to *P*, even if there are abnormal circumstances in which *R* would have occurred while *P* was not instantiated. Still, normalcy is notoriously hard to specify, and there may be no way of specifying it which marks a determinate threshold in the probability distribution, below which *R* is not sensitive to *P.* In that event, there will be vagueness in what counts as being sensitive to a property. There will be cases in which it is neither determinately the case that *R* is sensitive to *P*, nor determinately the case that *R* is insensitive to *P.*

However, I proposed *Sensitivity* as a necessary, not sufficient condition on visual representation of a property. Accordingly, vagueness in a representation's sensitivity might or might not entail semantic vagueness in that representation. Even if *Sensitivity* is a vague condition, it might be determinate that visual experience *represents* some precise determinable property, to which (it's determinate that) the experience is sensitive. This depends on further constraints on visual content, which it is not my project here to explore. If Searle is right that it is always determinate what conscious experience represents, then what visual experience represents is to some extent arbitrary, from the point of view of *Sensitivity.* But that is no objection to taking *Sensitivity* to be among the

conditions necessary for visual representation of a property. If *Vague Indeterminacy* is objectionable, we need not embrace it.

A different proposal about vagueness in visual representation is as follows:

Vague Determinacy

> It is indeterminate which of several maximally determinate properties a visual experience represents.

On this approach, it is determinate that a visual experience represents *some* maximally determinate property within the range to which the experience is sensitive, but it is indeterminate which of these maximally determinate properties the experience represents. Where the object of experience has one of these properties, it is then indeterminate whether the experience represents things as they are. *Vague Determinacy* is an alternative to *Indeterminacy*, since according to *Vague Determinacy* visual experience does represent some property to which it's not sensitive. By contrast, according to *Indeterminacy* visual experience represents a determinable property to which it is sensitive; there is no indeterminacy about which more determinate property is represented, since more determinate properties are not represented at all.

To the extent to which Searle's intuition is compelling, we have a reason to reject *Vague Determinacy* which is not a reason to reject *Indeterminacy*: according to *Vague Determinacy*, it is indeterminate what the content of conscious experience is. However, I adopt a more conservative strategy against *Vague Determinacy*. To motivate *Vague Determinacy*, we would need a good reason for saying that visual representation outstrips *Sensitivity*. I am arguing that visual phenomenology, at least, does not provide such a reason.

Finally, we should distinguish between the visual representation of determinable properties and blurry vision. The visual representation of determinable properties is a form of inexact vision. Michael Tye identifies blurry vision with inexact spatial vision:

> In these cases, one simply loses information. … In seeing blurrily, one undergoes sensory representations which fail to specify just where the boundaries and contours lie. Some information that was present with eyes

> focused in now missing. In particular, there is less definite information about surface depth, orientation, contours, etc.
>
> Tye 2002: 147-8

Now it is true that blurry vision 'loses information' in this way. But blurry vision does not consist only in this loss of information. Blurry vision is not *identical* with inexact spatial vision. Consider the case of seeing a distant object. Here vision loses 'definite information about surface depth, orientation, contours, etc', as compared to when you see the object close up. But distance vision need not be blurry. I take it that this is grounds for rejecting Tye's account of blurry vision, rather than grounds for thinking that visual experience represents equally definite spatial information about an object, however far away the object is.[32] Visual experience of determinable properties need not be blurry.

Once we distinguish between *Indeterminacy* and these other, related claims about visual experience, it is hard to see what phenomenological objection to *Indeterminacy* there could be. Where visual experience represents determinable properties, the experience may have the phenomenology characteristic of experience of a clear, precise, unified shape, colour and so on. I see no phenomenological grounds for insisting, in addition, that visual experience always presents us with maximally determinate properties – for example that is presents us with spatial properties defined in terms of spatial points.

[32] Citing Tye's discussion, Block (2010: 51-2) argues that determinable properties cannot figure in the phenomenology or ordinary, non-blurry visual experience. Block does not consider the case of distance vision.

Chapter 3

'INATTENTIONAL BLINDNESS' & PARTIAL REPORT [33]

3.1 Attending & Seeing: A Constitutive Connection?

To attend to something, in the sense which this dissertation is about, is to be conscious of it in a certain way. To see something is also to be conscious of it in a certain way – to be visually conscious of it. What's more, you often attend to the very things you see. For example, you're currently not only seeing the words on this page but also attending to them, focussing your consciousness on them in order to read.

What, then, is the relationship between attention and visual experience? It's clear that attention affects what we see. For example, if something takes your interest you can direct your gaze at it, and thereby come to see its details. But is there also a closer connection? In particular, is attention constitutive of visual experience?

In this chapter I will explore the connection between visual experience and *covert* conscious attention – conscious attention which involves no observable behaviour such as movements of the head or eye. I ask whether this form of attention is constitutive of visual experience. Covert attention is often described as operating like a spotlight, concentrating awareness and information-processing where its beam falls (e.g. Posner, 1980). We might spell out this analogy in either of the following ways:

Independence

> Attention ranges like a spotlight over the objects you see in the environment around you, concentrating consciousness on them. In that sense, attention illuminates objects whose character is independent of it. Likewise, attention ranges over the contents of visual experience, illuminating episodes of conscious vision without constituting them.

[33] Much of the material in this chapter appeared, in a slightly different form, in Stazicker 2011.

Dependence

> Attention ranges like a spotlight over the objects you see in the environment around you, concentrating consciousness on them – in that sense illuminating them. But visual experience is not always fixed independently of attention, there in advance to be illuminated by it, so to speak. Rather, attention is sometimes one of the factors constituting an episode of conscious vision and the way it presents things.

Of course, we could define a notion of visual experience such that *Dependence* is true. We could choose to treat visual experience as an episode which includes the conscious effects of attention. But that would avoid the interesting question about the structure of experience. It would invite the response that there is a further conscious visual episode, constituted independently of attention, which is only one element in visual experience as we're understanding it. A more promising approach is to tackle the question head-on, asking whether there is in fact always some further conscious visual episode constituted independently of attention.

What is it for one conscious episode to constitute another, or for one episode to depend constitutively on another? Sometimes philosophers use the term 'constitutive' to mean *essential*. For example it's argued that being of chemical composition H2O is constitutive of, or essential to, being water. But the term is also used to express a weaker connection – weaker in the sense that it's entailed by but doesn't entail essential connection. Scientists and philosophers investigating the neural basis of consciousness sometimes use 'constitutive' to describe a *minimal sufficient condition*: the constitutive basis of a particular conscious episode is the set of factors jointly sufficient for the phenomenal character and representational content of the conscious episode, and such that no proper subset of this set of factors is so sufficient. For example, following Cristof Koch, Ned Block tries to identify the 'minimal set of neuronal events and mechanisms jointly sufficient for a specific conscious percept' (Koch 2004: 16; Block 2005: 46). Block equates this set of events and mechanisms with the 'minimal neural basis of the phenomenal *content* of an experience, that which differs between the experience as of red and the experience as of green'.

This constitutive connection between physical and conscious phenomena may also entail constitutive connections *among* conscious

episodes: if the physical basis minimally sufficient for a conscious episode *A* includes the physical basis minimally sufficient for a conscious episode *B*, then *B* is itself part of the minimal sufficient condition on *A*; *B* is among the set of factors sufficient for the character and content of *A*, and such that no proper subset of this set of factors is so sufficient; *A* depends constitutively on *B.* I'll pursue the question whether, in this sense, visual experience sometimes depends constitutively on covert conscious attention.

To see the sense in which a minimal sufficient condition is a constitutive condition, consider a mountain that's partly made of – or partly constituted by – some granite. Being made of granite is not essential to being a mountain. Plenty of mountains have no granite in them. Nevertheless, the granite is constitutive of this particular mountain in an interesting sense:[34] the granite is among a set of factors jointly sufficient for the appearance, structure and causal powers of the mountain, and no proper subset of this set of factors is so sufficient. This same notion can be applied to events and episodes as well as to objects. The 1906 eruption of Vesuvius was constituted in part by an expulsion of magma from some rock, in that the expulsion of magma was among a set of factors jointly sufficient for the appearance, structure and causal powers of the eruption, and such that no proper subset of this set of factors was so sufficient.[35]

It is controversial exactly how sufficiency should be understood here. Konrad Marc-Wogau (1962) introduced the notion of a minimal sufficient condition into the philosophical literature to capture the historian's idea of a cause. Koch and Block have in mind a modally stronger condition, a condition on a conscious episode which can be contrasted with its mere causes. The simplest proposal would be that *A* is minimally sufficient for *B* only if it's absolutely impossible for *A* to occur without *B* occurring. But it is controversial whether physical factors are sufficient in this way for conscious episodes. Perhaps minimal sufficiency can be understood in a modally weaker way which nonetheless preserves the distinction between constitutive or minimal sufficient conditions on

[34] Or more than one interesting sense: the granite is a spatial part of the mountain; this connection may entail the further, more abstract constitutive connection I elaborate. I focus on the more abstract connection, because it applies more obviously to conscious episodes.

[35] We should accept this constitutive connection whether or not we also accept, with Davidson, 1969, that the 1906 eruption of Vesuvius is *identical* with a complex event of which the expulsion of magma is a part.

the one hand, and mere causes on the other. Specifying such a condition would be a substantial project, and I don't want to get involved in that here. I'll take *Dependence* to claim, and *Independence* to deny, that conscious attention is part of a minimal sufficient condition on some episodes of conscious vision, where the sufficiency here is modally as strong as the sufficiency connecting physical factors with conscious episodes – however strong exactly that turns out to be.

I think progress here requires that we give centre-stage to the indeterminacy of visual experience. In this chapter, I will show how underestimating visual indeterminacy lends spurious support to *Independence*. Then, in Chapter 4, I will show how taking visual indeterminacy seriously makes it plausible that *Dependence* is correct.

It is often suggested that experiments of two kinds bear on the question whether *Dependence* or *Independence* is correct. *'Inattentional blindness'* experiments are said to support *Dependence*, while *partial report* experiments are said to support *Independence.* In both cases, I think the suggestion is a mistake. In fact, once we appreciate the indeterminacy of visual experience, we can see that the mistakes here are quite symmetrical: in both cases, evidence about attention is not adequately dissociated from evidence about visual experience.

In §3.2, I discuss an argument for *Dependence* from the 'inattentional blindness' experiments. In §3.3 I discuss an argument for *Independence* from the partial report experiments.

3.2 'Inattentional Blindness'

Arien Mack and Irvin Rock (1998) argue for an extreme version of *Dependence*: the striking claim that 'attention is essential to conscious perception'; we do not consciously experience unattended objects of perceptual processing. They call this 'inattentional blindness'.[36]

[36] My criticism could easily be adapted to apply to other influential experimental work on 'inattentional blindness' – for example Simons and Chabris' (1999) argument that we are blind to unattended dynamic events.

Figure 11

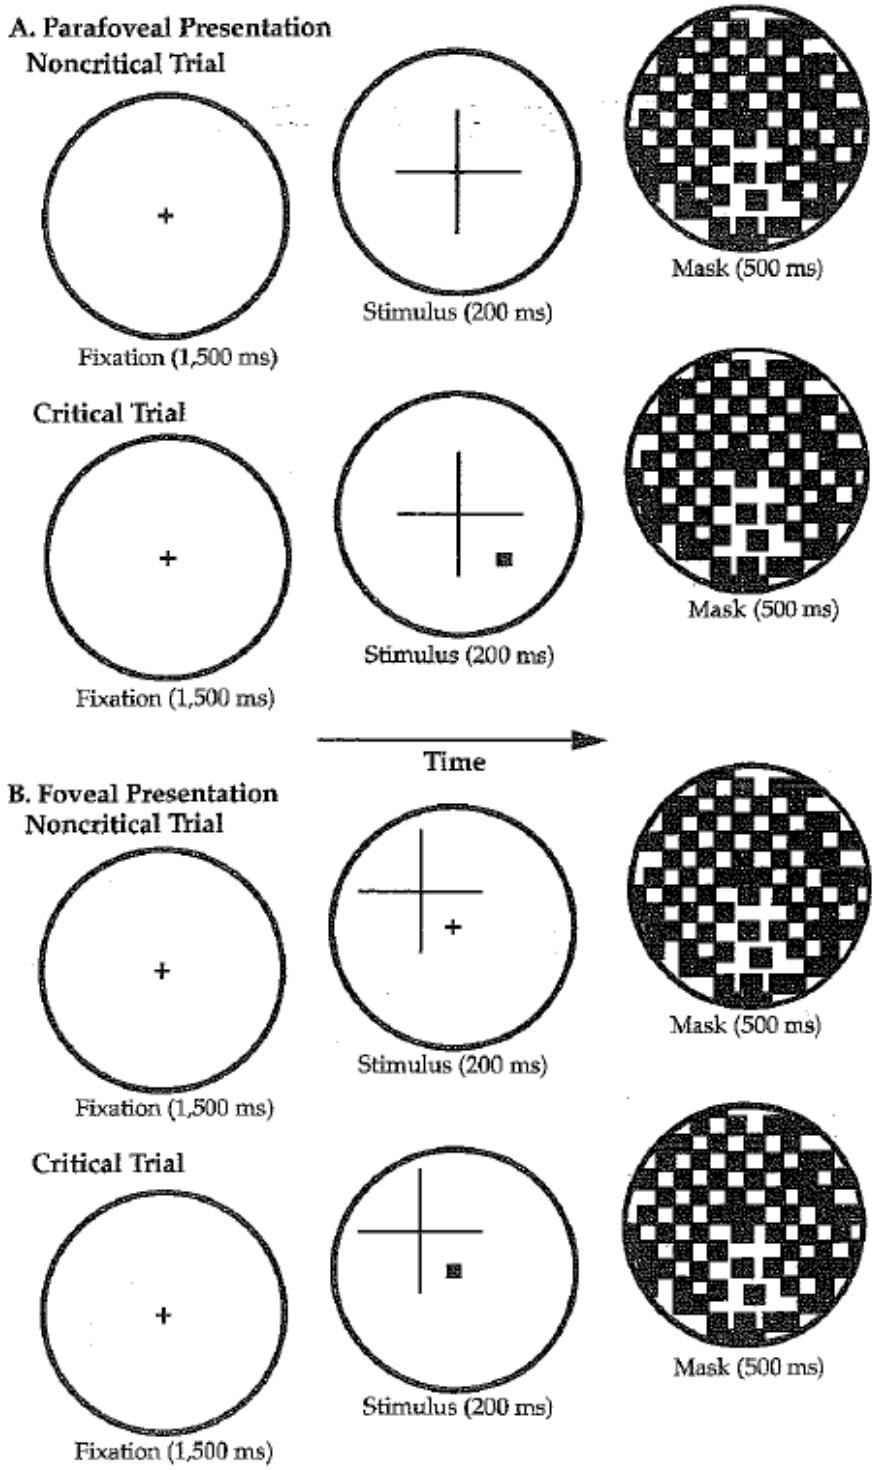

Reproduced from Mack and Rock 1998: 57, by permission of Arien Mack and Oxford University Press, Inc. (copyright 1998).

Mack and Rock's experimental method is illustrated in Figure 11. Subjects were given a task to occupy their attention. They had to report the longer arm of a cross. (Note that what is manipulated here is exactly the familiar sort of conscious attention involved in selecting something you see for thought.) The cross was presented for 0.2 seconds either at fixation or peripherally within 2.3° of fixation.[37] Fixation was held constant while covert attention was manipulated in this way. This process was repeated in four trials, except that in either the third or the fourth trial a further stimulus accompanied the cross, without warning. This stimulus was located in a quadrant of the cross, either at fixation or within 2.3° of fixation, depending on the position of the cross. On this critical trial, once subjects had reported the longer arm of the cross, they were immediately asked also to report whether they had seen anything that had not been present on previous trials, and if so to identify it or select it from an array of four to six objects. The idea was that these latter reports would reveal what subjects were visually conscious of, under a condition of inattention. Various stimuli were used, to test visual experience of various phenomena under a condition of inattention. These stimuli were all clearly visible when subjects attended to them: in a full-attention control trial, practically no subjects failed to report the additional stimulus.

On average, 25% of subjects reported that they had seen nothing other than the cross in the critical trial. More surprisingly, when the additional stimulus was at fixation, between 60% and 80% of subjects reported that they had seen nothing other than the cross. To explain this, Mack and Rock suggest that attention was actively inhibited at fixation. Since it is difficult to direct your visual attention away from fixation, tasks which require you to do this inhibit attention at fixation. To support their interpretation, Mack and Rock used an amended procedure designed to inhibit attention to peripheral locations. On the trials prior to the critical trial, attention was attracted to a peripheral location, by a stimulus irrelevant to the task. This seems to have led to inhibition of attention to the location of the irrelevant stimulus: 50% of subjects now reported having seen nothing other than the cross on the critical trial with the additional stimulus at this peripheral location.

---

37 Recall that the point of fixation is the point at which vision is literally most focussed. It's dictated by the relative positions of the eye and external stimuli: the point centred on the fovea, at a certain distance, is the point of fixation. You can attend covertly away from fixation (Posner 1980).

Still, these experiments might suggest that you can be visually conscious of something without attending to it in up to 75% of cases, depending on the circumstances. However, Mack and Rock adduce evidence for the view that some residual attention was responsible for visual experience in these cases. They reduced the size of the cross presented in their experiments, so that the additional stimulus was presented well outside the quadrants of the cross. With this amendment in place, subjects reported that they had seen nothing other than the cross on 66% of trials. This suggests that, in the earlier procedures, attention to the cross or to its location could spread to the additional stimulus; once the additional stimulus was far enough from the attended area, subjects reported seeing it far less often.

Mack and Rock claim that, if something were to fall both outside the region to which a subject was primarily attending, and in a region in which attention was actively inhibited, the subject would report not having seen anything there in almost 100% of cases. They argue that when something is completely unattended, the subject wholly lacks visual experience of it.

Perhaps one could object that Mack and Rock's data don't support a conclusion as strong as theirs. But I want to focus on a deeper problem with their approach. Where the subject reports seeing nothing other than the cross, Mack and Rock take this to show that she was not visually conscious of anything other than the cross. They assume that reports to the effect that nothing was seen are definitive evidence that nothing figured in visual experience. Why should we assume that? The problem is especially pressing given that we're dealing with a question about attention. Reporting on what we see requires us to attend to it. So the failure to report an object of visual experience might reflect a failure to attend to the object, rather than an absence of visual experience of the object. Suppose you want to defend *Independence.* Then you should be unmoved by Mack and Rock's argument. They beg the question against your view, by assuming that no conscious experience of an object is in place, where no attentive mechanisms of report on the object are in place.

The problem here runs really quite deep. There is no obvious way to test whether *Dependence* or *Independence* is correct, because in potentially probative cases reports and other fairly explicit decisions form our only convincing evidence concerning the character and presence of conscious experience. We are not limited to asking subjects to describe their experiences. For example, Mack and Rock did not only ask their

subjects whether they had seen anything other than the cross. They also asked them to identify what they had seen from among an array of stimuli, or – if they reported having seen nothing – to guess which of these stimuli had been presented. In general, we can ask subjects to make decisions which reveal what they experience consciously, rather than simply asking them to describe their experiences. But the evidence from such decisions is as problematic as the evidence from fully explicit verbal reports. Mack and Rock found that subjects who reported seeing nothing were also unable to identify the stimulus which had been presented. But this failure might reflect either (i) that the subjects were not visually conscious of the stimulus, or (ii) that, though subjects were conscious of the stimulus, they did not attend to it in the way required for this consciousness to form the basis for a reliable decision. To assume that (i) is the correct interpretation is to beg the question.

On the other hand, there is no obvious way to argue for interpretation (ii) either, because without reports or fairly explicit decisions we lack compelling evidence for the presence of consciousness. We might try looking beyond such decisions and reports for behaviour which reveals conscious awareness. But other sorts of behaviour are not sufficient to distinguish between the presence of conscious awareness and the presence non-conscious information-processing. For example, Mack and Rock found that the subjects who reported seeing nothing other than the cross were nevertheless affected by the additional stimulus, in ways which showed up in their behaviour. Where the additional stimulus was a word (e.g. 'Flake'), these subjects were significantly more likely than control subjects to use this word to complete a relevant word-stem (e.g. 'Fla-') in a subsequent trial. This is not good evidence for consciousness without attention. The effect may result, as Mack and Rock suggest, from deep but non-conscious processing of unattended stimuli. In general terms, the trouble is this: the only behaviour distinctive of visual experience of a visible phenomenon ϕ consists in reports and other fairly explicit judgements about ϕ; this behaviour requires attention.[38]

[38] Of course, there is room for sceptical doubt even where someone exhibits this sort of behaviour. In principle, perhaps someone might behave in the relevant ways without being visually conscious of anything. But we should set aside this sort of scepticism for practical purposes. It is not because of sceptical doubts that other sorts of behaviour are inadequate to prove that someone is conscious. They are inadequate because we're dealing with difficult cases. For example, we do in general take someone's avoiding obstacles in the street as evidence that she is visually conscious of those obstacles, even

If only we could isolate the neural activity correlated uniquely with visual experience, we could establish whether this activity was present in the absence of processes which are correlated uniquely with attention. But of course discovering the neural correlate of visual experience would require us first to have established when visual experience occurs, and so to have established whether it occurs without attention.

Hilary Putnam (1981) is so impressed by the difficulties here that he thinks there is, in principle, no hope of discovering an answer to the question whether there is unreportable consciousness, and no hope of discovering which brain activity is correlated uniquely with visual experience. Because verbal reports are indispensable as evidence of visual experience, Putnam claims that no explanatory grounds could justify a conclusion as to whether the mechanisms of report must be in place in order for someone to be visually conscious. This, he says, is 'a case to be *legislated* rather than fought over' (p.92); there is no fact of the matter, currently eluding discovery by us, about whether you could be conscious yet unable to report it; rather, we can choose which way we want to talk about consciousness.

Scientists working in this area sometimes seem to accept something like Putnam's anti-realist proposal. Victor Lamme (2005) says that we should identify visual experience with recurrent processing in areas of the brain which process visual stimuli. Recurrent processing is a distinctive kind of neural activity. It can be distinguished from feedforward processing, which moves roughly in one direction from the retina through a hierarchy of visual areas. Recurrent processing, by contrast, involves interactions between areas reached later in the feedforward sweep and areas reached earlier. Recurrent processing can also be distinguished from the processing distinctive of attention: recurrent processing does not entail the kind of neural memory involved in the selection of one stimulus at the expense of others.

In this paper, Lamme offers no reason for thinking that he has *discovered* an identity between visual experience and recurrent visual processing. He cites evidence for the view that recurrent processing, rather than mere feedforward processing, is necessary for visual

---

though a sceptic might doubt this. The question of what exactly enters her consciousness under conditions of inattention is more difficult, not because scepticism here is more pressing, but because the question requires us to tease apart different cognitive phenomena which are related in complex ways.

experience. For example, anaesthesis suppresses recurrent processing but not feedforward processing. But he gives no argument for the view that recurrent processing is sufficient for visual experience. Rather, he suggests that we 'redefine' visual experience as recurrent processing, since recurrent processing is a neurophysiologically identifiable phenomenon which admits of empirical investigation. Thus he legislates against the proposal that visual experience depends constitutively on the mechanisms of attention, in addition to recurrent processing.

The best way to resist this sort of anti-realism would be to find a way to argue for, rather than stipulate, a proposal about the relationship between attention and visual experience. In §3.3, I will criticise an attempt to argue that visual experience is independent of one sort of attention. Then in Chapter 4 I will argue that visual experience does sometimes depend constitutively on attention.

3.3 Partial Report

In a different paper, Lamme proposes the following account of the connection between attention and visual experience:

> Attention is a separate selection process, which is in principle independent of the conscious [visual] experience. Attention is a limited capacity, bottleneck-like, process, that allows stimuli to be processed deeper or faster, and which is necessary for storage in a durable working memory store or for a conscious report about stimuli.
>
> Lamme 2004: 863

He takes this account to be positively supported by some partial report experiments in which he was involved (Landman et al. 2003). So we have here an argument, rather than a stipulation, about the connection between attention and visual experience. But Lamme does not say very much about *why* he takes the partial report experiments to support this view. He just says that this is the 'logical solution' to interpreting the experiments.

However, Ned Block (2007) articulates the argument in detail. Block introduces a notion of 'cognitive access' which meets Lamme's description of attention, and argues by appeal to partial report experiments that visual experience is constitutively independent of

cognitive access. I should note that Block does not himself equate cognitive access with attention. So Block could take a more pluralist view of attention, along the lines which I will articulate in Chapter 4. Nonetheless, Block's argument can be treated as an argument for *Independence*, by those who identify attention with the phenomenon Lamme describes.

Block is acutely aware of the problem I raised in §3.2, and my discussion there owes much to his. He thinks we can use subjects' reports of their experiences, combined with their discriminatory capacities, to argue that visual experience is constitutively independent of cognitive access, in a way which avoids the problem I raised. He claims that subjects in partial report experiments report visual phenomenology which outstrips what they can access cognitively; their reports indicate that they were conscious of details in a stimulus, though they could not say what those details were. He claims that these subjects' performance in discrimination tests provides evidence that their reports are accurate. On this basis, he argues that 'phenomenology and cognitive access are based at least partly in different systems with different properties' (2007: 494).

By itself, this does not show that visual experience or phenomenology is constitutively independent of cognitive access. To complete the argument, Block turns to neurological evidence. Distinct 'coalitions' of neural activation compete for dominance both in visual areas of the brain at the back of the head and in frontal areas associated with access for reports. The coalitions in frontal areas are triggered by dominant coalitions in the back, but losing coalitions in the back may still be strongly activated (Kouider et al., 2007; Deheane et al., 2006). Block argues that this provides a mechanism apt to explain the 'overflow' of visual experience beyond what subjects access: while frontal activation is required for cognitive access, strong coalitions in the back of the head are sufficient for visual experience, whether or not they are dominant; thus visual experience or phenomenology has a greater capacity than cognitive access. If this proposal about the mechanisms of visual experience is correct, then the mechanisms of cognitive access are not 'a necessary part of a neural sufficient condition' on visual experience (489); neither these mechanisms of access nor the conscious episodes they constitute are among the factors constitutive of visual experience; visual experience is constitutively independent of this sort of attention.

I think Block's interpretation of the partial report experiments underestimates the indeterminacy of visual experience. Once we take

proper account of visual indeterminacy, we can see that Block's argument runs into trouble in much the way I described in §3.2: the only visual phenomenology for which Sperling's experiments provide evidence may, for all the experiments show, be accessed; so Block does not succeed in demonstrating that phenomenology overflows accessibility. And without the premise about 'overflow', his argument cannot get off the ground.

As evidence for his claims about 'overflow', Block cites the partial report experiments of George Sperling (1960). Sperling's first experiment was as follows. Subjects were presented for half a second with a grid of between six and twelve letters, with either three or four letters in each row. (See Figure 12 – or you can look at a live version Block's website.[39]) Subjects had to maintain fixation at a point in the centre of the grid. When asked to report which letters the grid contained, each subject could accurately report between 3.8 and 5.2 letters on average. These results were unaffected by the size of the grid, so subjects could accurately report fewer than half of the letters in the largest grids. Yet these subjects 'insist that they have seen more than they can … report afterwards' (Sperling 1960: 1); they 'said that they could see all or almost all of the letters' (Block 2007: 487).

Figure 12

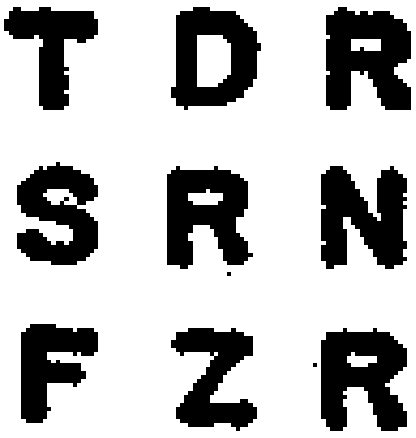

[39] http://www.nyu.edu/gsas/dept/philo/faculty/block/demos/Sperling320msec.mov. The presentation here is slightly shorted than Sperling's, so the task may be a little harder.

In subsequent experiments, trained subjects heard a tone immediately after the visual stimulus ceased. The pitch of the tone indicated which row of the grid they should report – a high tone for the top row, a low tone for the bottom, or a middle tone for the middle row. They were able to report which letters were in the indicated row with a high degree of accuracy – each subject getting between 75% and 100% right after practice, even when faced with the largest grid. This demonstrates that subjects processed information about the specific shape of almost every letter in the grid, even though they could access much less information when their attention was not cued. They must have processed information about the specific shape of almost every letter in the grid, because even once the stimulus was gone they could access that information if appropriately cued.

Block makes a stronger claim. He claims that Sperling's subjects enjoyed conscious, phenomenal experience as of the specific shapes of all or almost all the letters. He claims that this conscious perceptual content was inaccessible, in the sense that subjects were unable to report the specific shapes or identities of nearly half the letters. In that sense, he says, visual experience 'overflows' accessibility.[40]

In what sense were subjects visually conscious of the specific shapes of the letters, according to Block? Well first of all they saw them, and drew on their memory of what they saw in their reports. Onset of the auditory cue was controlled by the same switch which turned off the visual stimulus. So the cue was heard sufficiently soon after the visual

[40] Block also cites the partial report experiments by Landman et al. (2003). These combine a Sperling-style paradigm with a 'change blindness' task: when uncued, subjects could not accurately compare the size and orientation of some of the rectangles in two serially presented grids; when they were cued to the position of a rectangle in each grid, with the cue appearing on a grey screen between the presentations of rectangles, their performance improved. Block takes this to show that subjects were conscious of some rectangles' sizes and orientations, prior to and independently of accessing them. Despite the name, what this task tests for is not awareness of a visible change, but rather a capacity to notice a *difference* – a complex fact about the stimuli, rather than a feature of any individual stimulus (cf. Dretske 2004). This introduces a complication I don't want to discuss in detail: on the face of it, the data are compatible with uncued subjects having had cognitive access to the sizes and orientations of all the rectangles in both stimuli, and lacked only the capacity to compare them. In any case, my criticism of Block's discussion of the Sperling experiments applies equally to his discussion of the Landman et al. experiments.

stimulus ceased for it to be plausible that subjects were still seeing the recently-departed stimulus when they heard the cue and attended to the indicated row of letters. But Block notes that subjects in experiments similar to Sperling's sometimes report deliberately sustaining 'visual imagery' of the stimulus, after strictly speaking they have ceased to see it. He suggests that, if subjects are not still seeing the stimulus when they hear the cue, they draw on 'visual imagery' that's sustained after the stimulus has disappeared. And he counts this as a kind of visual experience (2007: 532).

Block is going on two sets of data: subjects' performance in the task of identifying letters in cued rows, and their verbal reports. The first set of data demonstrates that information about almost all the letters' specific shapes was processed independently of subjects' access to it, but it does not demonstrate that the subjects were *conscious* of these specific shapes independently of access. As we saw in Part 2, evidence that information was processed in your visual system is not necessarily evidence that you were conscious of that information. Block argues that the relevant information was processed in higher-level vision, rather than merely processed in early vision or simply recorded at the retina. But that does not demonstrate that subjects were conscious of this information. Precisely what's at issue here is whether higher-level vision is sufficient for visual experience, independently of the mechanisms of cognitive access.

A great deal turns, then, on Sperling's subjects' reports of their experiences. Unfortunately we don't have any quantitative data about these reports. But let's accept Block's claim that subjects reported seeing all or almost all of the letters. One option here would be to deny that we can take these reports at face value. That is what I proposed concerning the reports of Mack and Rock's subjects. But this case differs significantly from that one. Mack and Rock's subjects said that they had seen nothing other than the cross. Presumably, they reflected on their visual experience and found nothing there other than the cross. This admits of two explanations of roughly similar simplicity: (i) they were visually conscious of nothing other than the cross; (ii) though they were visually conscious of the additional stimulus, they failed to notice this. These explanations are of roughly similar simplicity, because (ii) requires us to accept only that noticing what we see is mediated by limited-capacity attentional mechanisms. The case we're dealing with now is different. The idea would be that Sperling's subjects did not really (consciously) see all or almost all of the letters; rather they thought they saw that many, but in fact they saw

only a maximum of 5.2 on average. If Sperling's subjects are wrong about their visual experiences they must have made a positive error about them, finding something there which was not in fact there, rather than simply failing to find something which was there. This requires a more elaborate explanation.

Some authors do suggest that Sperling's subjects might think they see more than they actually see (Dehaene et al. 2006). And others suggest a way of explaining how we might come to think of ourselves as visually aware of every part of an array in front of us, even though we're visually aware of only the portions of the array to which we attend (O'Regan and Noë 2001). The explanation is that, in order to discover whether you're seeing something, you attend to it, thereby becoming aware of it; this can give rise to the illusion that you were aware of the thing all along. The trouble with this line of response to Block is that it presupposes an interpretation of the 'inattentional blindness' experiments (Block 2007: 493). Only someone already convinced that visual experience requires attention would take such experiments as evidence that we systematically think we see things which we don't in fact see. We can't just make that assumption against Block's argument to the contrary. And without that assumption there's no reason to prefer this more complicated explanation to the simpler explanation that takes Sperling's subjects' reports at face value.

So let's take their reports at face value: Sperling's subjects could see all or almost all of the letters. Now Block takes this to show that their conscious visual experience contained more specific detail than they could access:

> [A]lthough one can distinctly see all of the 9-12 objects in an array, the processes that allow one to ... identify the specific shapes are limited by the capacity of 'working memory,' allowing reports of only about 4 of them.
>
> Block 2007: 487

But why should we accept that the information of which these subjects were conscious outstrips the information they were able to report? As I have said, their performance at the discrimination task demonstrates that they processed information about the specific shape of almost every letter. But why should we take it that they were conscious of this information?

Given *Spatial Indeterminacy*, we know that subjects experienced determinable shapes of the letters, rather than their maximally determinate shapes. So an alternative proposal is that subjects were conscious of all the letters, just as they said, but that they were not conscious of any shapes more determinate than the shapes they could access cognitively. Rather, they were conscious of some of the letters in a less determinate way, a way which matches the specificity of the reports they could offer. If they were like me they could report nine characters arranged in a grid, and they could report the identities of only about three or four of those characters. Similarly, the alternative proposal asserts, their visual experience should be characterized as an experience as of nine characters arranged in a grid, only three or four of them appearing sufficiently determinately to appear as a specific letter.

On this alternative proposal, the effect of cuing was not to allow access to a subset of the information of which Sperling's subjects were conscious. Rather, the effect was to alter their conscious experience such that some information became more determinate in it. This effect might have occurred either in a continuing episode of conscious seeing, or in an episode of sustained conscious visual imagery. Note that this interpretation doesn't require any *retroactive* effect of attention on conscious experience; attention made information more determinate in whatever episode of visual experience was ongoing at the time. According to this alternative interpretation, visual experience did not outstrip cognitive access. Rather, subjects were conscious of the items in the array with a degree of determinacy which precisely matched the determinacy of their cognitive access. If this alternative interpretation is as plausible as Block's, his argument for 'overflow' fails.

Block's interpretation of the experiments is of course consistent with *Indeterminacy*. *Indeterminacy* by itself does not tell us how determinate the shapes were, which subjects experienced the letters as having. However, there are sound empirical reasons to think the alternative relevant and plausible. Consider the crowding effect again. This is in the first instance an effect on subjects' capacity to identify alphanumeric characters: when a letter is presented in the periphery of a subject's visual field, surrounded by other letters, the subject is sometimes unable to identify the letter (Bouma 1970; try it for yourself with Figure 13). The effect is more marked where letters are presented fairly briefly, as Sperling's letters were presented, and in many instances the effect is mitigated by attention: attention to the crowded letter enables a subject to

identify it, just as attention to a letter enabled Sperling's subjects to identify it. As I said before, the mechanisms of crowding lie deep in the visual cortex: these are limitations on the determinacy with which we see, not limitations on cognitive access to a constant visual array (Levi 2008).

Figure 13

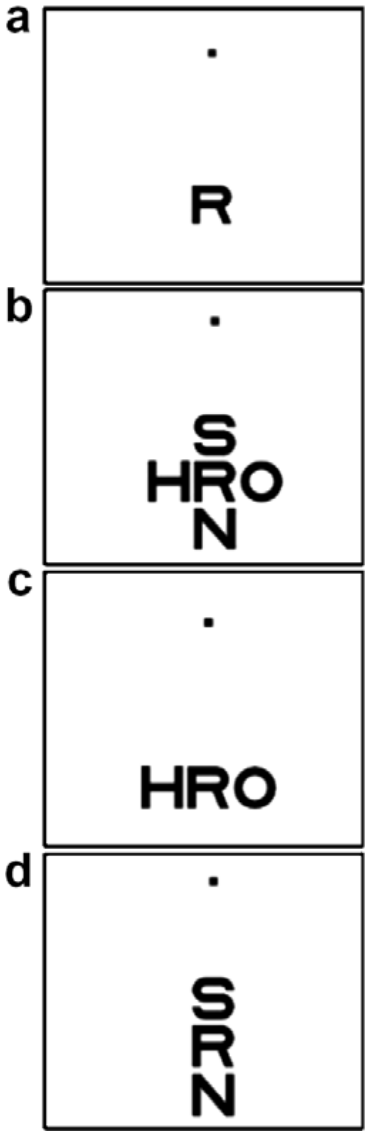

Fixate the dot in box (a), and try to identify the single letter below. Then do the same for central letter of the group, in boxes (b), (c) and (d).

Reproduced from Levi 2008: 636, by permission of the author and Elsevier (copyright 2008).

To be clear, I am not arguing specifically that Sperling's subjects suffered the crowding effect. For one thing, crowding typically occurs only for peripheral stimuli, and some of the letters which Sperling's subjects couldn't identify may have been at foveal locations. Rather, studies of the crowding effect provide evidence that subjects do, under some circumstances, experience letters insufficiently determinately to be able to identify them; what's more, in these cases, subjects come to experience letters more determinately, and are able to identify them, when they attend to them. This makes my alternative interpretation of the experiments empirically relevant.

Furthermore, quite aside from the crowding effect, we know that attention increases the resolution of visual processing. In at least some cases, attending to an item allows you to experience more determinate properties of it than you can experience when not attending. (I will explain the evidence for this in detail in Chapter 4.) This evidence likewise makes it an empirically relevant possibility that subjects in Sperling's experiments came to experience a shape sufficiently determinate for identification of a letter only when they attended to that letter. If that is the correct interpretation, then visual experience did not 'overflow' cognitive access in this respect.[41]

---

[41] I have emphasised the empirical reasons for taking this alternative seriously. But it is also striking that the philosophical tradition which emphasises perceptual indeterminacy is bound up with a view of the relationship between consciousness and attention that's diametrically opposed to Block's. For example, Leibniz held that ordinary perceptual consciousness consists in apperception of your own internal representational states. By 'apperception', Leibniz means 'attention to what is in us' (1714/1989, §4), and he works with a notion of attention similar to Block's notion of cognitive access: to attend to something is to concentrate higher cognition on it, or on a representation of it, either because of your antecedent interests or because it grabs attention exogenously (1704/1981, §54). So for Leibniz the machinery of consciousness is identical with the machinery of something like cognitive access. Now Leibniz also held that the determinacy of perceptual consciousness varies with apperception. For example, when you attend to the sound of the sea, your overall experience is made up of 'minute perceptions' of the sounds of individual waves; because these minute perceptions are not individually attended, they 'constitute that *je ne sais quoi*, those flavours, those images of sensible qualities, vivid in the aggregate but confused as to the parts' (1704/1981, §54). For Leibniz, confused representations can be contrasted not only with vivid representations, but also with clear representations, and a clear representation is defined as a representation which is sufficient for one to identify what's represented (1704/1981, II.xxix). So each part of the sound of the sea is experienced, but experienced in a way

Block suggests that cognitive access might represent Sperling's letters only 'under a general concept like "alphanumeric character"' (p.487). My objection is that visual experience might have a parallel indeterminate character, so that there is no mismatch between the content of visual experience and the content of cognitive access. In his published replies to criticism of his paper, Block draws a distinction between 'generic phenomenology' and 'specific phenomenology'. He applies this distinction as follows:

> For the Sperling experiment, the relevant generic/specific difference would be that between a phenomenal presentation *that there is* an array of alphanumeric characters and a phenomenal presentation of specific shapes of all or most of the items in the array. I argued that … there was specific phenomenology involving all or almost all of the items, as well as generic phenomenology.
>
> Block 2007: 531

Block marshals several arguments for his claim that subjects enjoyed 'specific phenomenology'. I address the key points below.

Block cites subjects' reports that they experience all or almost all of the items in the array. Such reports do not support the conclusion that subjects enjoy 'a phenomenal presentation of specific shapes of all or most of the items in the array'. They don't support this conclusion, because you may be conscious of an item, and indeed conscious of the item's shape, without being conscious of its specific shape. You might be conscious of a letter 'F', say, as having a certain general sort of shape, without being conscious of it as an 'F'. This involves more than Block's 'generic phenomenology' – more, that is, than being aware that there is an array of alphanumeric characters – because it involves awareness of the particular item. But it involves less than Block's 'specific phenomenology' – less, that is, than being aware of the specific shape of the particular item you see. Block attempts to divide conscious visual content into two categories, the specific and the generic. In this way, he artificially excludes the possibility of visual content with a degree of specificity that falls between his two options. He artificially excludes *Indeterminacy*, which I have argued we must accept. And the artificially excluded contents fit subjects' reports of

---

insufficiently determinate for one to identify it, because the overall sound is the object of attention. We might compare the individual letters in Sperling's grid.

the Sperling experiments just as well as Block's 'specific phenomenology' fits those reports. So Block takes the experiments to favour his view that visual experience is independent of cognitive access only because he casts the possibility of indeterminate visual content in terms of a false choice.

This same false choice is central to some of Block's other arguments that Sperling's subjects enjoy 'specific phenomenology'. Block cites evidence that subjects can report accurately whether they are conscious of a whole array, or conscious only of part of it. Subjects report seeing a whole 12-square grid if and only if they are also able accurately to report which of the squares have dots in them (Loftus & Irwin 1998; Brockmole et al. 2002). This suggests that we should take at face value subjects' reports that they are conscious of all or most of the items in Sperling's array. And it suggests that Sperling's subjects enjoy phenomenology more specific than Block's 'generic phenomenology', given that generic phenomenology doesn't suffice for awareness of any individual item in the array. But it does not suggest that Sperling's subjects enjoy 'specific phenomenology' of all or almost all of the items in the array. It doesn't suggest this, because it's silent on the question of whether subjects are conscious of the specific shapes of the items in the array.

Block notes that 'subjects report no … phenomenological shift' when they shift their attention to the specific items on which they report. Now my alternative interpretation of the Sperling experiments does require that subjects' phenomenology shift, in the sense that they become conscious of more specific details of the items to which they shift their attention. But it is not at all obvious that this shift is something subjects would, or could, report. Suppose you experience some of the items in Sperling's array with a shifting degree of determinacy. Suppose that you're first conscious of a letter-like shape in the bottom corner of the grid, and that when you shift your attention to that shape, you become conscious of the shape as an 'F'. How could you distinguish this from a case in which you were conscious of the shape as an 'F' all along, though you weren't attending to this aspect of it? Attention to what you see is required if you're to report on changes in the determinacy of your visual experience. But if attention to what you see *effects* changes in the determinacy of your visual experience, keeping track of the changes will be difficult at best.

Block takes it that the shift in phenomenology should be introspectively obvious, because he takes the shift to be between 'generic phenomenology' and 'specific phenomenology'. That sort of shift might

well be obvious, given that it would be a shift between, first, content which might be captured in a proposition about the array which makes reference to no specific item, and subsequently, content which might be captured in a detailed picture of all the items in the array. But the choice between 'generic phenomenology' and 'specific phenomenology' is false, and the shift need not be so dramatic.

Block also appeals to a different sort of report about experiences of the Sperling experiment. Subjects claim, he says, to be 'reading their answers off of the visual impression that was in existence before the cue'. Here again, we have no formal data about exactly what subjects say. But let's accept that they say something along these lines. Note that subjects do *not* say that they had a detailed visual impression which was unaltered by their shifting attention. To be sure, they were visually conscious of the whole grid before their attention was cued. And to the extent to which talk of visual impressions is appropriate here, subjects read their answers off their visual impressions. But this does not entail that their visual impressions were unaltered by the acts of attention involved in reading off their answers. Perhaps subjects' impressions became determinate with respect to specific shapes only where they cognitively accessed those specific shapes. If I ask you to form a visual image of your kitchen, and then ask you to report the specific shape of your stove, the image you enjoy will change as you shift your attention to the stove. The image will change in that, *inter alia*, you will represent the stove more determinately. Still, you were 'reading your answer off of the visual impression'.

The idea that Sperling's subjects read their answers off an image does not support Block's claim that there was 'specific phenomenology' prior to cognitive access to specific shapes. Nor does anything else about the Sperling experiment support the claim that visual experience 'overflows' cognitive access, rather than precisely matching the content of cognitive access. I have not ruled out the possibility that subjects were conscious of details they were unable to report. But there is equally no reason to think that they were so conscious. Where reports are non-specific, it may be that consciousness matches them. So neither set of experiments favours Block's claim that consciousness overflows cognitive access, any more than its denial. It is only because Block underestimates visual indeterminacy that he seems to avoid the problems described in §4.2. The problem facing his interpretation is quite symmetrical with the

problem facing those who appeal to 'inattentional blindness' experiments to defend *Dependence*.[42]

3.4 Making Progress

The foregoing leaves us with the *impasse* I identified in §3.2. Reports and other fairly explicit behaviour form the only good evidence concerning visual experience. For this evidence to be present, attention must be present. So we cannot argue for *Dependence* by appeal to the absence of this evidence, on pain of begging the question. Conversely, Block cannot establish that things figure in visual experience independently of whether you access them cognitively, since reports demonstrate that a feature of the stimulus figured in visual experience only where, or to the extent to which, the subject attended to that feature.

To escape this *impasse*, I think we need to note two points. First, the debate about *Dependence* and *Independence* which I have described is framed in terms of a certain specific sort of attention – the sort of attention in play when you focus your thoughts on something in the way required for verbal report. (This is clearest in Lamme's discussion, but it is also implicit in Mack and Rock's methodology.) Focussing your thoughts in that way does, I take it, entail attention. But it is not at all clear that attention occurs *only* when you do this. In Chapter 4, I will argue for a more liberal analysis of attention, according to which attention is on some

[42] Van Gulick (in Block 2007) suggests, as I have been, that visual experience might be indeterminate and so match cognitive access in its content. But Van Gulick doesn't spell out indeterminacy as I have (Grush's criticism is probably closest to mine), and Block's response is independent of his specific/generic distinction: Block says there's 'experimental evidence to the contrary'; 'cognitive access seems to be more of a binary phenomenon than Van Gulick supposes'. The evidence Block cites does not tell against the possibility I have raised – that cognitive access matched visual experience, in that both visual experience and cognitive access represented some items in Sperling's array without representing them as specific alphanumeric characters. This evidence concerns the 'attentional blink' (Sergent & Dehaene 2004). Subjects are asked to identify two sets of letters (e.g. XOOX and FIVE) from among a series of brief four-letter presentations. Given certain delays between the target presentations (especially 0.3 seconds), subjects fail to report the second target. When they're asked to rate, on a scale from 1 to 10, how visible the second target was, they rarely choose intermediate values. This suggests that whether or not the attentional blink prevents access to a particular letter is an all-or-nothing matter. But *that* might be an all-or-nothing matter, even if the shapes cognitive access represents vary in determinacy.

occasions a distinctively visual phenomenon, with its neural basis in the visual system, rather than in cognitive processes or the processes of working memory.

Second, the debate about *Dependence* and *Independence* which I have described is framed in terms of a very demanding criterion for constitutive dependence. Mack and Rock try to show that attention is necessary for visual experience – which entails that attention is part of a minimal sufficient condition on *every* episode of visual experience. Block tries to show that cognitive access to details of a scene is not necessary for visual experience of them; then, by appeal to his hypothesis about separate neural mechanisms, he argues that cognitive access is *never* part of a minimal sufficient condition on visual experience. These theses are certainly interesting. Perhaps they hold the promise of telling us whether or not visual experience essentially involves attending or accessing. But we can settle the constitutive question about *Dependence* and *Independence* without arguing for anything so bold. Minimal sufficient conditions on a token episode of conscious vision need not be conditions on *every* episode of conscious vision. In Chapter 4 I will argue that, in some token cases, your attending is part of a minimal sufficient condition on your seeing as you do. That is, attention is sometimes constitutive of visual experience; or *Dependence*.

The indeterminacy of visual experience is central to the positive argument of Chapter 4, just as it was central to the negative argument of §3.3. I will argue that, in some token cases, conscious attention takes the form of an increased determinacy in visual experience. In cases of this kind, the episode of attention is constitutive of the episode of visual experience (and *vice versa*). By way of introduction to this claim, consider some of Maurice Merleau-Ponty's comments about attention. Merleau-Ponty railed against an intellectualist account of attention, and he connected his criticism with his conviction that visual experience is indeterminate. Here is his critical characterisation of the intellectualist account:

> The function which reveals [the contents of visual experience], as a searchlight shows up objects pre-existing in the darkness, is called attention. … It is the same in all acts of attention, just as the searchlight's beam is the same whatever landscape is illuminated. … Inattentive perception contains nothing more than and indeed nothing other than the attentive kind. ... We are not called upon to analyse the act of attention as

> a passage from indistinctness to clarity, because the indistinctness is not there. Consciousness does not begin to exist until it sets limits to an object.
>
> Merleau-Ponty, 1945/1962: 30-32

Merleau-Ponty complains that, as the intellectualist conceives it, attention is exclusively post-perceptual: attention ranges over contents of visual experience which are fixed independently of attention. Hence the analogy with a searchlight, which ranges over a landscape the contents of which are fixed independently of the searchlight. The problem with this account, he says, is that attention is in fact 'a passage from indistinctness to clarity' *within* the contents of visual experience – a way of making determinate, and setting limits to, what would otherwise not be clear. Hence there is no independently-fixed landscape over which the searchlight of attention roams.

Although Merleau-Ponty would not have thought of it this way, we can see this as an empirical dispute. As the intellectualist conceives it, to attend to something is to select it from among the independently-fixed contents of visual experience, as a focus of post-perceptual conscious awareness. As Merleau-Ponty conceives it, to attend to something can be to see the thing more determinately. Attention can be a distinctively modal phenomenon. I will argue that the empirical evidence strongly favours Merleau-Ponty in this dispute.

Chapter 4

THE ANALYSIS OF ATTENTION

4.1 James's Definition

> Every one knows what attention is. It is the taking possession by the mind, in clear and vivid form, of one out of what seem several simultaneously possible objects or trains of thought. Focalization, concentration, of consciousness are of its essence. It implies withdrawal from some things in order to deal effectively with others.
>
> James 1890/1950: 403

William James's definition of attention is often cited, but seldom taken seriously enough. In particular, the last couple of years have seen several attempts by philosophers to analyse attention; these philosophers quote James, but then go on to impose more specific necessary and sufficient conditions on attention. I will argue that this is a mistake. We should work with James's definition, because it accurately captures the phenomenon of conscious attention.

The recent philosophical analyses portray attention as essentially *amodal*, or requiring a form of awareness distinct from experience in any of the perceptual modalities, *disembodied*, or independent of bodily movement, and *agential*, or involving action on the part of the attending subject. Together, these commitments amount to a form of intellectualism about attention. Attention is treated as if it were a form of intellectual reflection, a deliberate operation of the mind independent of the body and perceptual senses.[43] None of these features is in fact essential to attention. Attention is in many cases a modal, embodied phenomenon,

[43] The view that consciousness involves action is more usually associated with *anti*-intellectualism (Noë 2004). However, in conjunction with the two further ideas criticised here, the idea that attention involves action ascribes to the attending subject a kind of control which is characteristic of idealised intellectual reflection. I should note that none of philosophers I discuss commits to *all three* of the ideas I criticise. Nonetheless, I think it's fair to say that the intellectualism I describe captures the zeitgeist of contemporary philosophy of attention.

constitutively independent of action on the part of its subject. James's definition allows us to recognise this fact.

Here I am concerned only with *conscious* attention – with the sort of attention which is familiar from our introspectible conscious lives. Scientists study a broader range of phenomena under the heading of 'attention', and some of them explicitly argue that there is non-conscious attention (Koch and Tsuchiya 2006). In limiting my focus to conscious attention, I leave it open what the relationship is between this and other phenomena to which the term 'attention' gets applied. It might be that conscious attention is one species of a broader genus, attention, other species of which need not be conscious. Alternatively, it might be that every non-conscious phenomenon is *ipso facto* altogether distinct in kind from conscious attention, related to it only by loose resemblance or linguistic accident.

Even when we limit ourselves to conscious attention, the question 'What is attention?' may not have just one correct answer. The same goes for the question (asked in a particular case) 'Does this count as attention?' These are questions about how some conscious phenomena fall naturally into classes – or perhaps, ultimately, just questions about how we should group these phenomena. Either way, there might be several answers which are equally good. Nonetheless, I argue that James's definition does better than its recent competitors in the philosophical literature. The argument lies partly in my explicit comments in this chapter, and partly in the theoretical work to which I put the definition in the chapters which follow.

As James says, we are all familiar with conscious attention. The notion figures in our everyday, folk-psychological understanding of ourselves and one another. You know, for example, that paying attention to what someone says is a good way to make sure you remember it. And you know how to manipulate someone else's attention, how to distract them or focus them on a task. In that spirit, James's definition is in the first instance a conceptual analysis, and I will defend it partly in those terms. But a good analysis of attention should also be consistent with the discoveries of cognitive science, as well as with the role of attention in successful philosophical theory. I will argue that James's definition meets these demands more fully than its recent competitors.

James's definition is successful partly because it does justice to both the commitments and the modesty of our everyday idea of attention. While it is sufficiently informative to capture the pre-scientific notion, it

leaves a certain amount open for empirical and theoretical discovery. We are told that attention consists in a 'focalization of consciousness', but the definition leaves it open what forms this focalization may take, and what psychological mechanisms make it possible.

As we saw in Chapter 1, James identified both perceptual and post-perceptual mechanisms as the mechanisms of attention. It turns out, I will argue, that this general view of the mechanisms of conscious attention is correct. Attention may be distinctively visual, distinctively post-perceptual or both, depending on the specific case. James's definition allows for this because it defines attention in terms of the 'focalization of consciousness' – in terms of dynamic and comparative facts about first-order episodes of consciousness, which may themselves be distinctively visual or distinctively post-perceptual. Indeed, I suspect it is only because attention consists in such facts that we can define attention at all. Precise analyses of psychological kinds are in general hard to come by: seeing, knowing and so on have proved remarkably resistant to analysis. Attention, however, may be analysed in terms of facts about these first-order psychological phenomena.[44]

To dispute the analysis of attention intelligibly, we need some minimal common ground. I assume that we can agree on some paradigm cases: looking at ϕ, listening to ϕ, noticing ϕ, inspecting ϕ and considering ϕ all entail attending to ϕ (White 1964). Attending in these ways sometimes involves quite subtle shifts of awareness. For example, Jeff Speaks (2010) discusses a case in which Figure 14 occupies the whole of your visual field. Without moving your eyes, you can shift conscious attention between, say, the intersection of the horizontal line and the second line from the left, and the intersection of the horizontal line and the second line from the right.

[44] Mole (2011) takes James to task for identifying attention with a disjunction of disparate physiological processes. According to Mole, only an adverbial theory of attention can provide an appropriately unified analysis (see §4.2 below); to treat attention as a process is to make a category mistake, and this is reflected in the diversity of the physiological processes associated with attention. While I agree that attention is multiply realised, Mole's criticism of James seems to me uncharitable, and his dismissal of the process view misplaced. Even if the 'intimate nature' of attention lies in disjoint physiological processes, unity is to be found in the 'focalization of consciousness' which these disjoint mechanisms realise. As we'll see, this focalization may be either a conscious process or a conscious state.

Figure 14

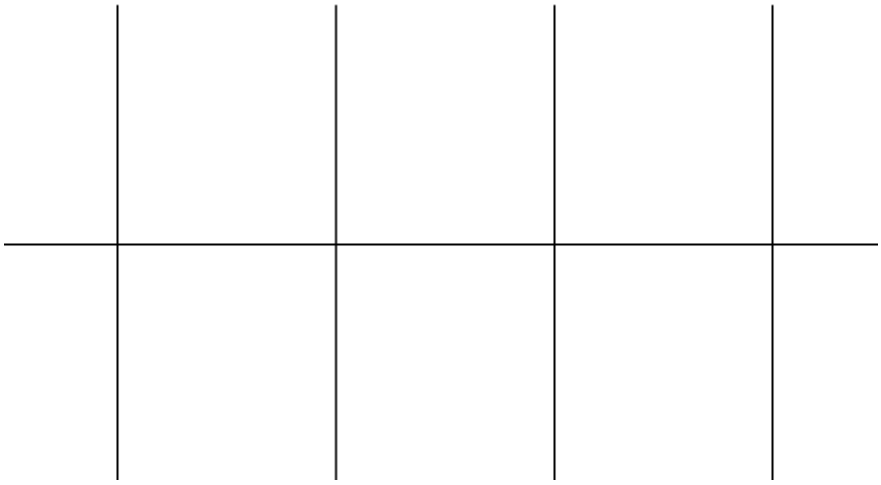

The philosophers whose analyses I will discuss all start with a case of this kind, noting its distinctive phenomenology, and then offer an analysis of the phenomenon of conscious attention which is exemplified there. Furthermore, they all agree with James that conscious attention to something is a way of selecting it as a focus of conscious experience. For example, in Speaks' case you alter your experience by taking first one intersection and then the other as a focus of conscious experience. In fact, every one of these authors quotes James, and none of them *disagrees* with his characterisation of conscious attention. Rather, they propose more specific necessary and sufficient conditions on attention, implying that James's definition is not specific enough.

Shortly I will show how the recent, more specific analyses get the phenomenon wrong, by casting as essential to conscious attention attributes which it does not in every case have. But first I want to show that James's definition actually gives us quite specific necessary and sufficient conditions on attention. We can set them out as the following three conditions.

Focalization

Necessarily, if *S* attends to ϕ, *S*'s consciousness is focalized on ϕ.

What is it for *S*'s consciousness to be focalized on ϕ? As James's gloss 'concentration' implies, *S*'s consciousness is focalized on ϕ only if *S* is *more* conscious of ϕ. Intuitively, when you attend to something, you are more conscious of it. But what is it to be more conscious of something? In what sense does consciousness admit of degrees?

One natural idea is that you may be conscious of something in more than one way, or in more than one mode of consciousness. As we'll see, recent philosophical work about attention focuses on the idea that, in addition to consciously perceiving an object, you may experience the object in a distinctively non-perceptual way. Different philosophers propose different accounts of this non-perceptual, attentive form of consciousness, but the general proposal is this: the focalization which James describes consists in your becoming conscious of an object through a non-perceptual mode of consciousness, in addition to being conscious of it perceptually. This kind of focalization involves a change or difference in the mode of consciousness by which you are aware of something, rather than a change or difference in what you are aware of. I'll call it *mode focalization*.

We can contrast this with *content focalization*, which involves a change or difference within the scope of a modality of consciousness. James says that when you attend to ϕ, ϕ is 'clear' to you. Traditionally, a clear representation is *inter alia* a representation which makes apparent the details of its object.[45] This suggests that consciousness also admits of degrees in the following sense: *S*'s consciousness may be focalized on ϕ in that more of the details of ϕ are apparent to *S*, within a mode of consciousness. For example, in §4.3 I will argue that your attending to ϕ sometimes consists in the fact that your visual experience represents more determinate properties of ϕ.

Just through introspection, it's hard to tell the difference between mode focalization and content focalization. When you attend to something you see, more determinate details of it become apparent to you. As we saw in Chapter 3, it is hard to tell whether this consists in your noticing details which you saw all along, or in those details becoming visible to you. The conscious phenomenon is familiar, but exactly what form it takes is not introspectively obvious. This suggests that James was

[45] See e.g. my discussion of Leibniz at §3.3, footnote 42. In addition, James's expression 'taking possession' might itself imply that ϕ must give up its details, so to speak, rather than keeping them from *S*.

right to characterise the familiar concept of attention in terms of the general notion of focalization; to disambiguate between mode and content focalization is to go beyond what's familiar from introspection. However, we can sometimes distinguish between these forms of focalization, using the techniques which vision scientists have developed to distinguish between visual effects on the one hand and effects on post-visual decision processes on the other. In §4.3, I will explain how the experimental evidence shows that there is content focalization in some paradigm cases of conscious attention.

Focalization is essentially contrastive: for your consciousness to be focalized on ϕ is for you to be more conscious of ϕ. The contrast may be with another object of consciousness, or it may be with consciousness of ϕ at a previous time. At a given moment you may be more conscious of ϕ than you are of χ. Equally, you may become more conscious of ϕ than you were previously. So attention may be either a state or a process, consisting in either a difference or a change in conscious experience. In many cases you attend in both these senses.[46]

Those contrasts are with actual facts about consciousness. James makes clear that attention also essentially involves a counterfactual contrast. Attention to something 'implies withdrawal from' 'what seem several simultaneously possible objects or trains of thought':

Selection

> Necessarily, if *S* attends to ϕ, *S*'s consciousness could have focalized on some distinct phenomenon χ, instead of ϕ. This possibility is introspectible by *S*.[47]

The counterfactual condition is not just that *S*'s consciousness might, in some possible scenario, have been concentrated on χ instead. For example, you experience more determinate spatial properties of nearby objects than

[46] Wu's (2011a) discussion of the distinction between synchronic and diachronic 'phenomenal salience' is helpful here.

[47] Assuming that *S* has the relevant concepts and is sufficiently self-aware. I am making some interpretive choices here. First, perhaps the appearance of alternative possibilities is part of the occurrent phenomenology of attention itself, rather than merely available upon introspection. The difference won't matter here. Second, perhaps it's necessary only that there *seem* to be alternative possible objects of attention – not that the possibility be genuine. I think the rest of the passage belies that interpretation.

you experience of far-off objects (§2.2). So had different objects been nearby, you would have experienced more determinate properties of them instead. It does not follow that you attend to everything nearby. Rather, the relevant counterfactual scenarios are ones in which *S*'s consciousness is concentrated on χ because *S* responds differently to a scene which is held fixed – because *S*'s consciousness *takes possession* of a different possible object or train of thought.

For example, attend again to one of the intersections in Figure 14. Here you could have attended elsewhere, in that you could have decided to direct your attention to a different intersection, or to something else altogether. In other cases attention is involuntary, and the relevant response is not an action or choice. For example, a sudden movement in peripheral vision sometimes attracts your attention. Suppose this happens at one of the intersections in Figure 14. The sudden movement causes you to respond in a way which focalizes your consciousness on the intersection, after the movement has occurred. In this case your consciousness could have focalized elsewhere, because there might have been sudden movement at a different intersection instead, or somewhere else altogether; as a result you might have responded differently to the same scene, after the movement occurred.

Finally, consciousness must be focalized on the object of attention 'in order to deal effectively with' it:

Purpose

> Necessarily, if *S* attends to ϕ, *S*'s consciousness is focalized on ϕ in order that *S*'s intentional states or processes be effective with respect to ϕ.

'Intentional states or processes' should be understood broadly, to include states of perceptual experience, thought and belief, as well as intentional actions. Similarly, the purposive 'in order to' should be understood as including not only the personal aims of the attending subject, but also the goals we may treat a perceptual system as having, when we assess how effectively it is functioning. So for example you might attend to a moving car, with the deliberate aim of avoiding it. But equally the car's movement might draw your attention involuntarily to it, independently of any personal aims you have. As we'll see, in the latter case attention takes the form of content focalization within visual experience. This focalization

itself constitutes the visual system's dealing effectively with the car, in that its details become visible to you. This is independent of your having any particular *reason* to want to see those details. Rather, we can treat your visual system as having the goal of revealing the details of an object, where that object is a potential threat (cf. Carrasco et al. 2008).[48]

4.2 Recent Analyses

The first standard contemporary claim that I want to dispute is:

Amodality

> Necessarily, if *S* attends to ϕ, *S* enjoys a non-perceptual mode of consciousness of ϕ.

The idea is that attention is a mode of consciousness over and above perceptual experience: in no context is experience of ϕ in one or more of the perceptual modalities sufficient for attention to ϕ; non-perceptual consciousness is also required. This is consistent, of course, with there being interesting connections between perceptual experience and attention. For example, *Amodality* is consistent with the idea that ϕ's being an object of perceptual experience causes or explains ϕ's being an object of attention. In computational terms, it is consistent with the idea that perceptual experience serves as the input for attention. But for someone committed to *Amodality*, what kind of psychological state or episode is conscious attention itself?

Speaks' suggestion here is in some ways the most straightforward. He proposes that we treat attention as an additional, *sui generis* modality of conscious experience: attention to ϕ is a distinctive way of being conscious of ϕ, analogous to but constitutively independent of seeing ϕ, hearing ϕ, or being conscious of ϕ through the other perceptual senses.

[48] Contrast Wayne Wu's (2011b) analysis of attention as 'selection for action'. According to Wu the action in question need not be bodily; thinking is a kind of action. Nonetheless, Wu imposes a more demanding condition on attention than James does, at least as I interpret James. Wu's condition is more demanding, in that perceiving something effectively – which James's definition allows as the purpose of attention – is not obviously a form of action (although see Noë 2004 for the view that perceiving is in general a form of action).

Consider again the two distinct experiences of Speaks' grid – attending first to one intersection and then to another. Speaks explains his proposal by way of a striking analogy:

> On this view, the two experiences of the grid could be compared with a pair of visual experiences identical except that one of them in accompanied by a toothache. While there is a clear difference in the total phenomenology of the subject during the two experiences, … there is … no difference in specifically visual phenomenology.
>
> Speaks 2010: 333

Of course, attention has a connection with perceptual experience which toothaches typically lack. Perceptual experience and attention often share an object, and Speaks notes that attention is 'constrained by' perceptual experience; the contents of attention are 'parasitic on' those of perceptual experience (2010: 334, 341). But this is a causal – perhaps computational – constraint on attention. It is consistent with Speaks' claim that attention to ϕ and perceptual experience of ϕ are distinct forms of consciousness, instances of which are constitutively independent of one another.

Speaks is concerned with the phenomenology of attention, rather than with its functional role in our mental lives. But on the face of it, attention to ϕ enables you to do and think things which you would not otherwise be able to do or think, with respect to ϕ. Declan Smithies develops an account of attention which has much in common with Speaks', but which connects the phenomenology of attention with its functional role. Smithies claims that 'attention is rational-access consciousness'. That is, to attend to ϕ is to access ϕ, or to select ϕ in consciousness, in a way which makes information about ϕ available 'for use in the rational control of action, reasoning and verbal report' (2011: 257). This control is rational in the sense that information about ϕ is 'accessible to the subject as a reason that justifies the subject in forming a belief or performing an action'. So for example attention to one of the intersections in Figure 14 allows you to respond, rationally, to a demand that you point to it or describe it as it appears in your visual experience.

According to Smithies, to attend to ϕ is to have this kind of rational access to information about ϕ. In his view, vision furnishes this kind of rational access to an object only where you enjoy visual experience of the

object. However, conscious attention itself 'is a distinctive mode of consciousness'. Like Speaks, Smithies says that conscious attention has a 'distinctive and proprietary phenomenology' which is 'arguably *sui generis*: why suppose that the phenomenology of attention can be reduced to the phenomenology of perception, action, cognition, or anything else?'

Someone committed to *Amodality* need not think of attention in this way. For example, Wayne Wu (2011a) argues that conscious attention to an object should be identified with conscious perceptual-demonstrative cognition about that object – roughly, with conscious thought which we can express by talking about 'that object', and which refers to the object in virtue of the way the object figures in perceptual experience. Here again, conscious attention to an object requires perceptual experience of that object. Thoughts of the kind in question are possible – in Wu's view – only where there is perceptual experience. But the thought itself forms a conscious episode over and above any episode of perceptual experience. In no context is any perceptual experience of ϕ sufficient for attention to ϕ. Attention is identified with a non-perceptual mode of consciousness.

The second standard claim I want to dispute is:

Agency

> Necessarily, if *S* attends to ϕ, *S* acts with respect to ϕ.

This commitment is most explicit in Sebastian Watzl's work. Watzl says that attention is an 'action' or 'activity', which 'consists in the conscious mental process of structuring one's stream of consciousness so that some parts of it are more central than others' (2011: 145). On the view Watzl favours, the conscious character of attention is a form of 'action-awareness', where the action in question is the action of attending of structuring your stream of consciousness.

Agency is also a clear consequence of other recent analyses. For example, Christopher Mole argues that 'attention is cognitive unison': to attend to an object *O* is to perform a task which involves *O*, in such a way that your performance of the task displays cognitive unison (2011: 73). Cognitive unison is defined as follows:

> Let α be an agent, let τ be some task that the agent is performing, and call the set of cognitive resources that α can, with understanding, bring to bear in the service of τ, τ's 'background set'.

> α's performance of τ displays cognitive unison if and only if the resources in τ's background set are not occupied with activity that does not serve τ.
>
> Mole 2011: 51

As attention is defined here, action with respect to an object is essential to attention to that object. The attending agent must be *performing* a task which involves the object, and she must be doing so in such a way that the resources which she can *with understanding bring to bear* on the task are directed only at that task. In short, the task must be the focus of her action.

The third standard commitment I want to dispute is:

Disembodiment

> Necessarily, if *S* attends to ϕ, *S* attends to ϕ by means of some action *A*, such that *A* is not a movement of *S*'s body.

The idea is that attending requires an action over and above any bodily movement: in no context is directing your perceptual organs to ϕ sufficient for attention to ϕ; a further, mental and non-bodily action is also required. For practical purposes, *Disembodiment* follows from the conjunction of *Amodality* and *Agency*. Movements of the body may bring about changes in the contents of perceptual experience, but according to *Amodality* these changes in perceptual experience never suffice for attending. And according to *Agency*, the further condition on attending is an action. It is true that we can use bodily movement as a means of carrying out non-perceptual cognition. For example, you can count on your fingers. However, bodily movement has no such role in attention. When you attend to something, the bodily movements which you carry out enable cognition by altering the contents of perceptual experience. For example, shifting your gaze brings new things into view. And according to *Amodality*, that does not suffice for attention.

For example, Watzl (2011) makes clear that the 'process of structuring one's stream of consciousness' does not consist in altering the contents of perceptual experience. In his view, the structure of the stream of consciousness consists in irreducible facts about what is central, and what peripheral, to conscious attention. So directing your gaze at ϕ never suffices for attention to ϕ. To attend to ϕ, you must in addition perform the mental and non-bodily action of making ϕ central in your stream of consciousness. Similarly, there is no sensory organ by means of which we

effect Speaks' *sui generis* sensory modality of attention. Rather, once all the facts about where you have directed your perceptual organs are settled, there is a further, non-bodily act to be carried out – the act of directing the *sui generis* modality of attention.

Consider again the psychologists' distinction between overt and covert attention. In the visual case, overt attention involves movement of the head or eye, while covert attention requires no such observable behaviour. Presumably Watzl, Speaks and others do not deny that overt attention occurs. Rather, the idea must be that overt attention to ϕ requires a further way of selecting ϕ as a focus of consciousness, in addition to any movements which direct your perceptual organs at ϕ.

Now what I want to dispute is just the claim that any of these features is essential to attention – that attention is necessarily amodal, agential or disembodied. I agree that attention sometimes takes the form of amodal conscious awareness, that it sometimes takes the form of an action, and that this action sometimes involves no bodily movement. However, attention need not have any of these characteristics.

I suspect that the contemporary philosophers' view of attention is informed by the mid-Twentieth Century psychologists' view of attention. As I said in Chapter 1, psychologists of that period hoped to discover a unique physiological process of attention; by and large, they believed in a distinctive, unitary 'bottleneck' process, by which some of the information registered by the senses is selected for further processing in limited-capacity mechanisms, while other information is discarded (Broadbent 1958).

This naturally suggests that conscious attention should be construed as a post-perceptual mode of consciousness, a way of selecting some of the contents of perceptual experience so that they may be further taken up in thought and belief (Figure 15). In recent years it has become clear that the mid-Twentieth Century psychologists' view was a mistake: a bewildering variety of physiological mechanisms realise the selective phenomena usually studied under the heading of attention; these mechanisms take effect throughout visual processing, rather than at a distinctive bottleneck location (Allport 1993). It does not immediately follow that the contemporary philosophers' idea about *conscious* attention is a mistake. It could, in principle, have turned out that only bottleneck-like mechanisms realise conscious attention, even though there are various other mechanisms of non-conscious attention. But that is not how it has turned out. Conscious attention may be a distinctively visual

phenomenon, realised by mechanisms early in the processing responsible for visual experience.

Figure 15

The mid-Twentieth Century psychologists' view of attention

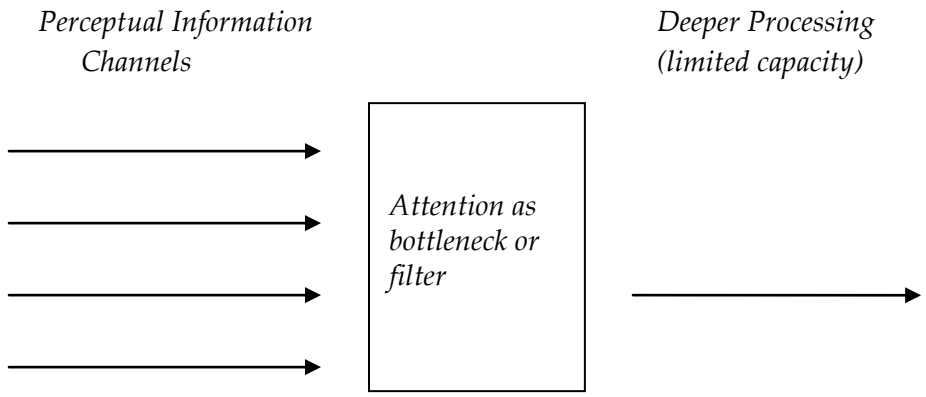

The current philosopher's view of attention

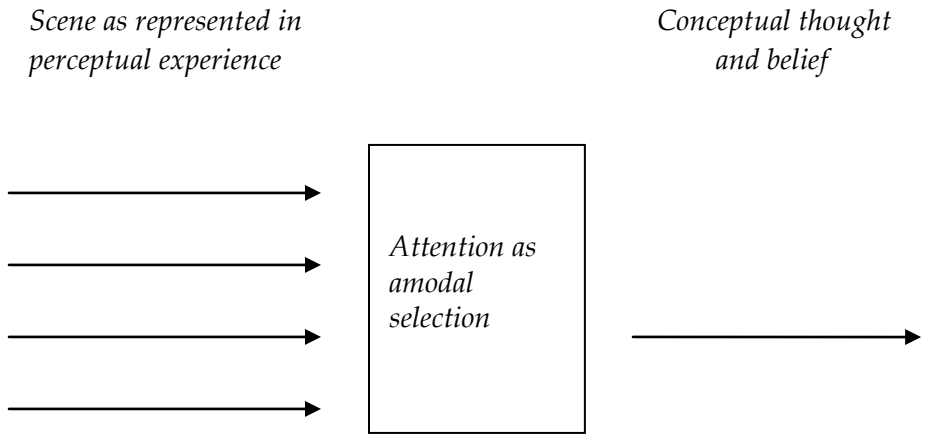

4.3 A Counterexample

Perhaps various different counterexamples could be adduced, to show that *Amodality*, *Agency* and *Disembodiment* are false. But I will explain how one body of work, and one set of experiments in particular,

provides a counterexample to all three theses. This work concerns the role of attention in visual spatial resolution.

Within covert attention, vision scientists further distinguish between voluntary and involuntary attention. Voluntary attention is directed by your aims and decisions, or by processes which, though not under your personal control, are nonetheless directed by endogenous, top-down controls – as with the automatic shifts of attention which occur when you enter a room. Involuntary attention, by contrast, is captured automatically and exogenously, by sudden changes in the visible environment.

It is a familiar fact that movement or change in peripheral vision draws your attention to its location. A light flashes, or someone fidgets during the talk, and you find your cognitive resources directed to the place where this happened. You might turn your head or eyes to this location, and you might actively concentrate on it. But whether or not you do these things, there is a short-term effect on your visual system: in an effect which peaks around 100 milliseconds after attention is cued and disappear after another 20 ms or so, processing in the visual cortex occurs faster for stimuli at the location in question, and contrast-sensitivity is increased for stimuli there – i.e. visual-cortical responses become more sensitive to differences in those stimuli (Carrasco and McElree 2001; Treue 2000). This is involuntary covert attention.

Involuntary covert attention occurs deep in the visual system, and it is automatic, rather than being the sort of thing that you can *do*. But this does not by itself tell us much about *Amodality* or *Agency*. So far we have no reason to think that these effects on the visual system amount to conscious attention – to something's being selected as a focus of conscious awareness. So far we just have some ways in which visual processing can be made more efficient with respect to an automatically selected target. However, Yaffa Yeshurun and Marisa Carrasco (1998) tested the role of involuntary covert attention in visual spatial resolution. And their results constitute a counterexample to both *Amodality* and *Agency*, because they show that involuntary covert attention does sometimes amount to something's being selected as a focus of conscious awareness – specifically, to its being selected as a focus of determinate visual experience.

Yeshurun and Carrasco used a texture segregation task. This is a task in which you have to report whether a stimulus contains a section in which lines have a unique orientation, relative to the lines in other

sections of the stimulus (Figure 16). In general, texture segregation becomes easier as the unique texture approaches fixation; the less peripheral the stimulus, the easier it is to tell whether it contains a unique texture. Holding fixation fixed, Yeshurun and Carrasco cued attention by flashing a small bar immediately above a location in a texture-segregation stimulus (Figure 17). When they cued attention to a peripheral location containing a texture-segregation target, they found that attending to that location made the task easier.

Figure 16

A texture-segregation stimulus

Figure 17

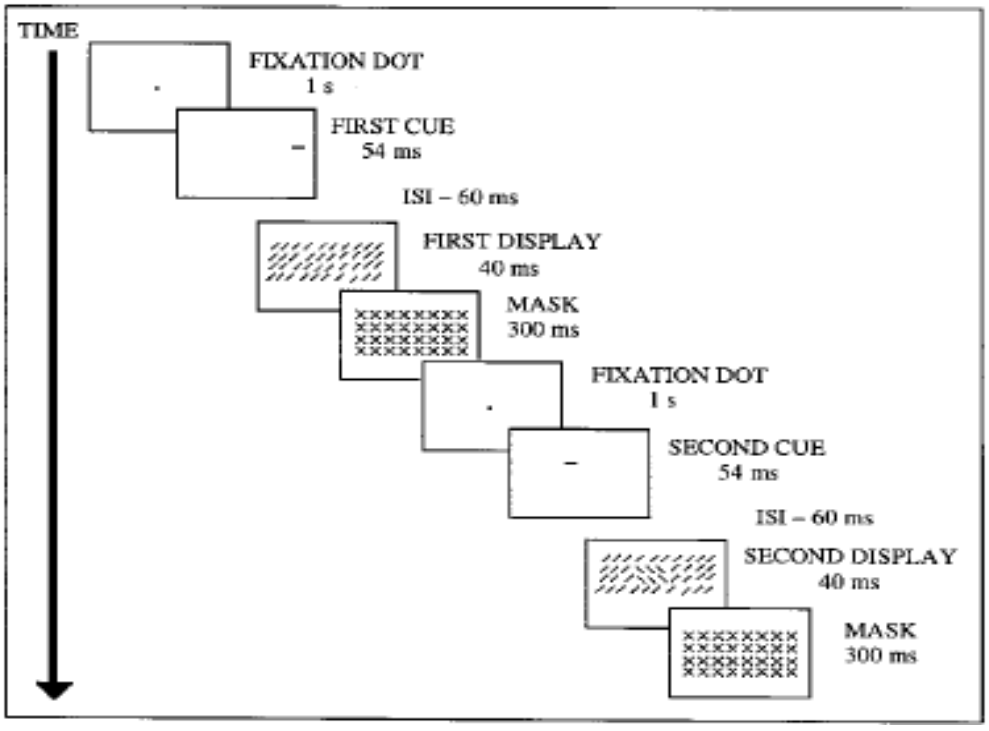

Reproduced from Yeshurun and Carrasco 1998: 72, by permission of Marisa Carrasco and Macmillan Publishers Ltd (copyright 1998).

This result is consistent with *Amodality*. In attending to the cued location, you might simply focus post-visual, amodal consciousness on it, and in that way become more aware of its distinctive features. Attending might in that way highlight a location so that your thoughts and decisions are concentrated on it, without affecting visual experience itself. The result is also consistent with *Agency*. Where conscious attention is concerned, at least, cuing to a location might just make that location salient, enabling you to engage in the act of attending to it.

However, Yeshurun and Carrasco devised a way of showing that the effect here is an effect on vision itself, rather than an effect on post-visual decision processes. Recall that visual spatial resolution is highest at the fovea, and progressively lower away from the centre of the eye. Early in post-retinal visual processing, information about the scene you see passes through varying spatial filters, attuned to various spatial frequencies. As the locations of these filters get further from the fovea, they are attuned to progressively lower spatial frequencies: higher-frequency detail goes unprocessed (DeValois and DeValois 1988). In general, seeing with higher spatial resolution is useful for texture segregation. It's because of this that texture segregation is in general easier near fixation. However, there are some texture segregation tasks in which seeing with too high a spatial resolution makes it harder to identify the unique texture. As Yeshurun and Carrasco put it:

> When a more global inspection of the display is required, for example when one is appreciating an impressionist painting, moving closer is not the optimal strategy
>
> Yeshurun and Carrasco 1998: 74

As a result, there are some texture segregation tasks at which performance drops as the unique texture approaches fixation. Yeshurun and Carrasco exploit this phenomenon to argue against the view that, in the cases they studied, attention merely affected subjects' decision-making procedures, leaving the visual signal static. They argue that attending enhances spatial resolution in the visual signal itself, by mimicking the effect of foveation, by effectively making the spatial filters for a location more finely tuned. In my terms – and assuming *Sensitivity* from Chapter 2 – they tested the hypothesis that visual experience was *not* static through changing

attention, that attention increased the determinacy of the spatial properties represented in visual experience itself.

In a task in which texture segregation is harder where spatial resolution is too high, they found that cuing attention to the unique texture *impaired* performance, when that texture was at or near foveation. This result cannot be explained in terms of the idea that attention simply improves your access to what's represented in a static visual experience. According to that proposal, attention should only improve the decisions you make, by devoting more decision-making resources to the relevant location. Attending to the unique texture should never diminish performance at texture segregation. To explain the data, we must accept that attention increases the spatial resolution of visual experience, replicating the effect of foveation.

From the point of view of an interest in conscious experience, Yeshurun and Carrasco's discovery is this: conscious attention – the selection of part of a scene as a focus of conscious awareness – sometimes takes the form of an automatic, selective increase in the determinacy of the spatial properties which visual experience represents that part of the scene as having. In Jamesian terms, the discovery is that the focalization of consciousness may take a visual form. More specifically, the effect which Yeshurun and Carrasco studied is a case of content focalization within visual consciousness. Subjects' visual experiences represented more determinate spatial properties of the cued stimulus than they did of the uncued stimulus. And their visual experiences represented more determinate spatial properties at part of the scene after it was cued, than they did before it was cued. So there was a focalization of visual consciousness, in both the *difference* and the *process* senses which I distinguished in §4.1.

This focalization is a modal phenomenon. It is a change in the properties which visual experience represents. And it does not require any action with respect to the attended object. Tuning of the spatial filters in early vision is not something that an agent can do for herself. If this interpretation of the experiments is correct, then both *Amodality* and *Agency* are false.

Furthermore, Yeshurun and Carrasco connect covert attention with overt attention, in a way which undermines *Disembodiment*. Their results provide evidence about the tuning of spatial filters in early vision, because involuntary covert attention replicates the effect of overt attention (i.e. of foveation): each form of attention hinders performance at the task.

Yeshurun and Carrasco explain this by hypothesising that each form of attention ensures that spatial filters sensitive to high spatial frequencies process information from the attended stimulus.

As I said, one might suppose that overt attention counts as attention only because – and to the extent to which – it is accompanied by an amodal focussing of consciousness; in addition to turning your eyes to the object of attention, you turn post-perceptual attention to it. That is one conception of what unifies overt and covert attention. Yeshurun and Carrasco's work suggests a different conception. Like attending covertly, overtly turning your eyes to part of a scene is itself a means of selecting that part of the scene as a focus of conscious awareness: by fixating part of a scene, you come to see it in higher resolution; you see more determinate properties of it. If this visual effect amounts to attention, then the means by which you attend overtly is itself a bodily action, and *Disembodiment* is false. From the point of view of an interest in conscious attention, covert involuntary attention consists in just the same conscious effect. The difference lies only in the means by which this effect is achieved: covert involuntary attention requires no bodily action.

There are two clear ways in which you might object, if you wanted to resist this package of claims about attention:

1. *Equivocation*

> If there is a sense in which automatic visual effects amount to attention, this is different from the sense of 'attention' at issue in contemporary philosophical analyses. So the effects are no counterexample to the analyses.

The idea is that involuntary covert attention is not a form of conscious attention in the sense in which philosophers are interested; in that sense, attention is always an amodal mental action.

Of course, you could just stipulate *Equivocation* with reference to your preferred analysis of attention. I don't deny that attention is in *some* instances an amodal, agential, disembodied phenomenon. So if you wanted, you could stipulate that you will use the term 'attention' to refer only to that phenomenon. But short of such a stipulation we have no reason to accept *Equivocation*.

One possible argument for *Equivocation* appeals to empirical work on the differences between voluntary and involuntary attention. Bill

Prinzmetal and Ayelet Landau (2008) argue that voluntary and involuntary attention are distinct processes, on the grounds that they involve distinct neural mechanisms and have distinct effects. Now there is an empirical dispute to be had here. Prinzmetal and Landau argue that involuntary attention affects visual processes only by making them faster; they do not consider Yeshurun and Carrasco's argument about the spatial filters in early vision. But there is also an underlying issue about where we should look for unity in an account of attention. For Prinzmetal and Landau, 'the critical question is whether these two forms of attention involve the same or different mechanisms' (2008: 43). By contrast, I have argued that unity is to be found at the level of consciousness: voluntary and involuntary attention are of the same psychological kind, in that they each involve a Jamesian focalization of consciousness; more specifically, voluntary overt and involuntary covert attention each involve an increase in the determinacy of the spatial properties which visual experience represents things as having.[49] I do not thereby imply that Prinzmetal and Landau say anything false. Rather, what they say about the various and distinct mechanisms of attention is compatible with my claim that these forms of attention are unified at the level of conscious experience.

I suppose someone hard-nosed might say that the mechanisms are all there really is, so really there is no unity here. Eliminativism of this kind is not well motivated in this case. Elimination of an apparent psychological phenomenon might be justified, if it turns out that this phenomenon would be inconsistent with what we know about the brain. For example Ramsey, Stich and Garon (1990) argued that there is really nothing which corresponds to the folk-psychological notion of belief, on the grounds that nothing in the brain could realise the causal role which folk psychology takes to be essential to belief. By contrast, it does not suffice for elimination that a high-level psychological phenomenon is realised by various different mechanisms. Compare the case of thinking. It turns out that some thoughts about an action like kicking involve effects in the motor cortex, while thoughts about the colours, for example, involve no such effects (Barsalou 1999). This hardly justifies a claim that there is really no such thing as thought, or that we equivocate when we say there are thoughts about both actions and colours.

[49] It is a difficult question, which I take up in Chapter 6, whether voluntary covert attention involves the same effect.

The contemporary philosophical analyses of conscious attention are, at least in part, analyses of a pre-theoretic, folk-psychological notion. Their proponents start with cases which we usually think of as paradigms of attention, and then offer analyses of the familiar form of conscious awareness that's distinctive of those cases. So a different argument for *Equivocation* appeals to the folk-psychological notion of attention, claiming that the automatic visual effect which Yeshurun and Carrasco discovered falls outside its extension.

There is, however, no good reason to accept this claim about the notion of attention. The folk-psychological notion to which the philosophers appeal concerns a phenomenal kind, a kind of conscious state which we take ourselves to enjoy in paradigm cases of attention. The case which Yeshurun and Carrasco studied is a paradigm of attention. That attention can be captured by a sudden peripheral movement is among the more obvious platitudes of folk psychology. Furthermore, as I have argued, the automatic visual effect which Yeshurun and Carrasco discovered shares a significant phenomenal kind with other effects which are paradigmatic of conscious attention. We can, through careful analysis and experiment, distinguish between different specifications of this phenomenal kind: mode focalization on the one hand, and content focalization within visual experience on the other. But why should we believe in a folk-psychological notion of attention which excludes the latter? On the face of it, we could not use such a notion competently in everyday first-personal or third-personal ascriptions of attention. The notion could be correctly applied only by someone able to distinguish a case in which visual experience remains fixed while further forms of cognition alter, from a case in which visual experience itself becomes more determinate. We saw in Chapter 3 that this distinction cannot be drawn either through first-personal introspection or through the means of third-personal interrogation usually available to most of us.

This is not to deny that attention often takes the form of mode focalization. In Chapter 6, I will suggest that some of the facts about conscious attention can be explained only in terms of mode focalization, and I will suggest that mode focalization should be understood as a cognitive phenomenon – a phenomenon of conscious thought. For now I just want to emphasise that, on the view defended here, conscious attention may take the form of either mode focalization or content focalization; indeed, I'll suggest that in many cases attention takes both these forms.

Admittedly, the relevant notion of attention should not be framed in exclusively phenomenological terms. Smithies introduces his discussion as follows:

> Ordinarily, we think of attention both in terms of its phenomenology and its functional role. Shifting one's attention from one thing to another affects one's overall phenomenology, but it also affects one's functional dispositions to think about or act upon the one thing rather than the other.
>
> Smithies 2011: 247

It's not obvious that seeing an object in more determinate detail should raise the probability that you will think about or act on that object, rather than other objects which you see less determinately. So *Equivocation* might be defended on these grounds: the content focalization which Yeshurun and Carrasco discovered lacks a functional role that's required by what 'we ordinarily think' of attention. However this defence of *Equivocation*, like the one above, employs an artificially restrictive notion of attention. Seeing an object in more determinate detail might not raise the probability that you will think about or act on that object, but it does affect *what* you are likely to think and do with respect to the object, as well as how you see it. Seeing in this way enables you to deal with the object effectively, in the sense spelled out in *Purpose* (§4.1). I know of no argument that the more restrictive notion is required by the folk psychology of attention. In fact, I suspect that the defunct 'bottleneck' hypothesis about the mechanism of attention is all that tempts us to the restriction, if anything does.

Structurally similar comments apply to the causes of attention, as opposed to its effects. To defend *Equivocation*, you might note that we ordinarily take attention to be something we do, not just something which happens to us. For example, we may hold someone responsible for failing to pay attention. However, it would be artificially restrictive to infer that attention is *only ever* something we do. In general, *X* may be something we do, while *X* also sometimes just happens to us. For example, falling is sometimes an action, something for which you may be held responsible. Think of the cheating soccer player who dives to the ground to get a free kick. But falling also often just happens to us. The soccer player might trip or be pushed. In fact some cases of falling are mixed: the soccer player might begin to fall passively, before engaging in a theatrical tumble. Or

consider weeping. Sometimes this just happens to us, but sometimes it is something we actively do. And some cases are mixed: once weeping comes over you, you can engage in it as an action. I propose that the same goes for attention: there are automatic visual effects which themselves amount to attention; once your attention is drawn to something in this way, you may further attend by actively concentrating on it.

A second way to resist my package of claims about attention is:

2. *Not Constitutive*

> Automatic visual effects may be correlated with attention, but attention does not consist in these effects.

The idea is that, although attention is accompanied by automatic visual effects, attention itself is essentially amodal and agential. For example Watzl (2011) argues that, although conscious attention causes you to see more determinate properties at the attended part of a scene, conscious attention does not consist in this visual effect.[50]

Watzl argues for *Not Constitutive* by appeal to:

Reproduction

> There are possible differences in the visible environment which reproduce, at an unattended part of the scene, just the increase in visual determinacy which attention brings about.

His argument is as follows. Take any scenario in which, because you attend to part *A* of a scene and not to part *B* of the same scene, you experience more determinate properties at part *A* than you experience at part *B*. Given *Reproduction*, this effect on visual experience may be reproduced in a case in which you are not attending to part *A*. Therefore,

[50] Like the study I cited (Yeshurun and Carrasco 1998), the study which Watzl cites (Carrasco et al. 2004) concerns involuntary attention. In cases of involuntary attention, it's really not plausible that amodal attention is a *cause* of increased visual determinacy. The mechanism of increased determinacy probably lies in early visual processes such as those to which Yeshurun and Carrasco appeal. However, Watzl's proposal could easily be amended to say that increased visual determinacy is merely correlated with conscious attention – the ground of the correlation being a common cause, namely the sudden movement or change which attracts attention.

your attending to part *A* does not consist in this effect on visual experience; the effect on visual experience must be merely an effect of attention, because it 'leaves out an essential aspect of the phenomenology of attention' (2011: 152).

For example Watzl says that, because the right side of Figure 18 is blurry, you see less determinate properties of it than you see of the left side. You do not thereby attend to the left side of the figure, or enjoy the phenomenology of attention with respect to it. Indeed you can attend to the right side of the figure instead, and it will still be blurry. Similarly, Watzl associates the determinacy of visual experience with the visible contrast of a scene. He takes evidence that subjects experience higher apparent contrast at the attended part of a scene (Carrasco et al. 2004) to be evidence that attention makes visual experience represent more determinate properties at that part of the scene. But, Watzl argues, this must be a mere effect of attention, since there could be a version of Figure 18 such that the right side is lower in contrast than the left – a version such that, in virtue of the properties of the figure itself, you see more determinate properties on the left side, even when you don't attend there.

Figure 18

From Watzl (2011), by permission of the author and OUP (copyright 2011).

I'd like to make two points in response to this argument. The first is phenomenological, the second logical. First, *Reproduction* is false. There are no possible differences in the environment which reproduce just the effect on the determinacy of visual experience which Yeshurun and Carrasco found. The differences in visual experience which Watzl considers are unlike those I explained above, in that they are not differences in visual

resolution. Measures of visual resolution abstract from changes in the environment. For example, spatial resolution is measured as the maximum spatial frequency to which a visual response is sensitive *as a function of stimulus contrast* (DeValois and DeValois 1988). As the contrast of a stimulus increases, so the maximum spatial frequency to which a visual response is sensitive must increase, if the spatial resolution of the response is to retain a constant spatial resolution. Therefore increased contrast at part *A* of a scene does not increase the spatial resolution of responses to part *A.* Similarly, to introduce blurriness into a scene is to change the spatial frequency of the stimulus. You do not change the spatial resolution of a response to a stimulus, by changing the spatial frequencies which are available for detection in the stimulus itself. (We should not confuse seeing in low resolution with seeing a low-resolution image such as a blurry photograph. Low-resolution images *are* typically blurry, but they may be seen in high resolution; the blur may itself be seen in determinate detail.)[51]

These are not just technical curiosities about how resolution is measured. They ensure that the phenomenology of increased visual spatial resolution is not reproducible through differences in the visible environment. Unlike differences in visual resolution, the differences in visual experience which Watzl considers are not *only* differences in the determinacy of the properties which visual experience represents. Rather, they introduce into the scene new properties which may themselves be represented at varying degrees of determinacy.

For example, the blurriness on the right of Figure 18 is itself a visible aspect of the scene – an aspect which has properties at various levels of determinacy, such that the determinacy of the properties you experience it as having may vary. As a result, it may not be true that you see more determinate properties on the left of Figure 18 than you see on

[51] As standardly measured, spatial resolution also abstracts from an object's proximity to the observer. Bringing an object closer to you allows you to see more determinate spatial properties of it, but this does not entail an increase in the spatial resolution of visual responses to the object. Spatial resolution is measured as the maximum spatial frequency to which a visual response is sensitive, where spatial frequency, in turn, is measured in cycles of change per degree of visual angle (§2.2). For example, 1 metre of a far-off part of a scene takes up fewer degrees of visual angle than 1 metre of a nearby part. So where spatial resolution in measured in this way, it is no higher for nearby parts of a scene than it is for far-off parts. We can't reproduce high-resolution experience of an object just by moving it close.

the right, even if you don't attend to the left. Perhaps you see more determinate properties of the dot on the left than you experience of the dot on the right. But the blur may itself be seen in determinate detail. So overall, you may experience more determinate properties on the right than you do on the left. For example, if your attention were cued to the right of the figure, you might experience the figure in that way.

Similarly, an increase in contrast on one side of the figure might lead you to experience *some* properties more determinately on that side of the figure. However, an increase in contrast at part of a scene is itself an aspect of the visible scene, an aspect which may be represented at varying levels of determinacy. (For example the brightness contrast which Carrasco et al. studied is usually measured as the difference between the maximum and minimum brightness of a stimulus, divided by its mean brightness.) So where part of a scene is relatively high in contrast, it is does not follow that, overall, you experience more determinate properties of that part of the scene. You might experience a less determinate contrast there.[52]

In this respect, differences in visual resolution are unlike differences in contrast and blurriness. Where visual experiences of parts *A* and *B* of a scene differ only in that *A* has a higher resolution, it *does* follow that, overall, you experience more determinate properties at part *A* than at part *B*. Moreover, suppose that there is some way of assessing the determinacy of the properties you experience, such that overall you do – as Watzl claims – experience more determinate properties on the left of Figure 18, even when you don't attend there. This does not suffice for *Reproduction*, since the determinacy effect is accompanied by a phenomenal change which is absent in cases of a straightforward increase in visual resolution: a change in the apparent blurriness or contrast of part of the scene. In this way, an increase in visual resolution is distinctive of attention. There are no possible changes in the environment which reproduce it. *Reproduction* is false, and the argument for *Not Constitutive* fails. To say that conscious attention sometimes consists in this visual effect is not to leave an essential aspect of the phenomenology of attention.

[52] Wu (2011) makes an argument similar to Watzl's, and similarly cites Carrasco's experiments about attention and contrast. I discuss these experiments in Chapter 5. They are less directly relevant than the experiments of Yeshurun and Carrasco 1998 to questions about the determinacy of visual experience.

Now for the second, logical point in response to Watzl's argument. Consider a phenomenal difference of the kind which Watzl focuses on, such as an increase in the apparent contrast of a scene. This does, I think, amount to a focalization of consciousness, and to a case of conscious attention. Where apparent contrast is increased, other aspects of a scene such as the shapes of things become more visible (Fuller and Carrasco 2006 – more on this in Chapter 5). Now unlike a straightforward increase in visual resolution, this phenomenal difference is reproducible in the absence of attention. But that doesn't show that something essential to the phenomenology of attention has been left out, when we treat it as a case of attention. Rather, there are further conditions on conscious attention, beyond its phenomenology. Recall that attention involves 'taking possession by the mind' of one object and 'withdrawal' from others (1890/1950: 403). *Selection* reflected this by requiring that the focalization of consciousness depend counterfactually on the subject's responses to a scene which is held fixed (§4.1). It's this condition on conscious attention which is failed, if we reproduce the increase in apparent contrast by altering the contrast of the scene, rather than by inducing a different subjective response to an environment which is held fixed.

Indeed, we can also make a parallel point about the condition I called *Purpose.* Consider an alternative to *Reproduction*:

*Reproduction**

> There are possible differences in visual processing which reproduce, at an unattended part of the scene, just the increase in visual determinacy which attention brings about.

*Reproduction** is surely true, even with respect to differences in visual resolution. Presumably there are possible differences in visual processing which would make you see some parts of a scene in higher resolution than others, in the absence of attention. For example, perhaps this could be achieved by manipulating the spatial filters in early vision. But this is no objection to my claim that attention sometimes consists in increased visual spatial resolution. Where attention consists in this, the further condition of *Purpose* must also be met: your consciousness must be focalized on ϕ in order that your intentional states and processes be effective with respect to ϕ. Where increased visual resolution fails this condition, it does not

amount to attention, even though the subject's experience has just the phenomenology of attention.

4.4 *Dependence* Again

An episode of attention may take the form of an episode of visual experience which meets the three Jamesian conditions of *Focalization*, *Selection* and *Purpose*. Because it meets these conditions, the episode of visual experience is constitutive of attention. In a case of this kind, attention depends constitutively on visual experience. What's more, and conversely, the episode of visual experience depends constitutively on the episode of attention, in the sense explained in Chapter 3. There I described two competing proposals about the connection between attention and visual experience:

Independence

> Attention ranges like a spotlight over the objects you see in the environment around you, concentrating consciousness on them. In that sense, attention illuminates objects whose character is independent of it. Likewise, attention ranges over the contents of visual experience, illuminating episodes of conscious vision without constituting them.

Dependence

> Attention ranges like a spotlight over the objects you see in the environment around you, concentrating consciousness on them – in that sense illuminating them. But visual experience is not always fixed independently of attention, there in advance to be illuminated by it, so to speak. Rather, attention is sometimes one of the factors constituting an episode of conscious vision and the way it presents things.

Dependence requires that in some token cases, your attending is part of a minimal sufficient condition on your seeing as you do: the episode of attention must be among the set of factors sufficient for the content and character of an episode of visual consciousness, and such that no proper subset of this set of factors is so sufficient.

Yeshurun and Carrasco's experiments show that visual experience sometimes meets the Jamesian conditions on attention; the 'focalization, concentration of consciousness' which James described sometimes takes the form of an increase in the determinacy of conscious vision. As a result, the episode of focalization or attention is among the set of factors sufficient for the changing content of visual consciousness, and such that no proper subset of this set is so sufficient. Someone attending in this way is, necessarily, also enjoying a visual experience with a certain changing content. And nothing short of attending in this way would suffice, under the circumstances, for enjoying this experience.

I suppose this is not the sort of answer which many philosophers and scientists expected, when they framed the constitutive question about attention and visual experience. For one thing, this answer tells us nothing about whether attention is essential to visual experience. But it is nonetheless an appropriate answer to the constitutive question, as the notion of constitution is usually understood by theorists who pose the question (e.g. Koch 2004; Block 2005): an episode of attention is part of a minimal sufficient condition on an episode of visual experience.

The constitutive question is often intended as a question about the neural bases of attention and visual experience. In cases of the kind Yeshurun and Carrasco studied, the constitutive dependence of visual experience on attention is reflected in the fact that the neural basis of visual experience includes the neural basis of attention.

Consider again the sort of proposal about the neural basis of visual experience which is usually associated with *Independence*. For example, Victor Lamme notes that processes in the visual cortex, in particular recurrent processes involving V4, are a good candidate for being among the set of factors constitutive of visual consciousness. Changes in these processes are implicated in changes in the course of conscious visual experience (Lamme 2005), which suggests that they form part of a minimal sufficient condition on that experience (Block 2005).

Now it is well-established that attention produces heightened selective responses in V4 (Moran and Desimone, 1985; Luck et al., 1997). These responses are a good candidate for being at least part of the neural basis minimally sufficient for the episodes of conscious attention on which I have been focusing – i.e. part of the neural basis minimally sufficient for the focalization of consciousness involved in seeing with a greater spatial determinacy (cf. Yeshurun and Carrasco: 74). For that same reason, they're a good candidate for being at least part of the neural basis minimally

sufficient for the changing content of visual experience. Whatever is minimally sufficient for the episode of conscious attention is also part of what's minimally sufficient for the changing episode of visual consciousness. So, where we're dealing with conscious attention at least, the same neural events sometimes realise or constitute both attention and visual experience.[53]

The constitutive connection is also significant from a phenomenological point of view. The connection may not be introspectively obvious. It may be hard to tell, just through introspection, what form the focalization of consciousness takes. Nonetheless, the constitutive connection is a feature of the phenomenology of attention. As Merleau-Ponty claimed, attention does not merely roam over contents of visual experience which are settled independently, 'as a searchlight shows up objects pre-existing in the darkness'; rather, attention marks a 'passage from indistinctness to clarity' within visual experience itself (1945/1962: 30-32).

Recent philosophical analyses imply a separation between visual experience and conscious attention which is neither neurologically nor phenomenologically apt. Conscious attention may take a visual form, such that we should understand conscious experience not just in terms of attention on the one hand and visual experience on the other, but also in terms of unified episodes of attentive visual experience. James's definition allows us to capture this fact about attention.

[53] This proposal is consistent with some recent claims about a neural basis of attention outside the visual cortex. For example, I don't dispute the suggestion of Moore (2006) that the *Frontal Eye Field* might drive involuntary attention. Rather, following James I treat 'attend' as a success verb: to attend is to achieve a focalization of consciousness. Accordingly, the neural basis of attention lies in the processes which realise this conscious change. Moore could agree, I take it, that these lie at least partly in the visual cortex. He reports that stimulating a monkey's *FEF* caused increased differential responses to differently oriented stimuli, in neurons in V4.

Chapter 5

VISUAL KNOWLEDGE, ATTENTION & PSYCHOPHYSICS

5.1 Two Roles for Attention

On the face of it, attention has an important epistemic role. We often explain why someone comes to know about something she sees, or why she does not come to know about it, by saying that she pays attention to it or that she does not pay attention to it. We take it that you come to know about the things you see by paying attention to them. My focus in this dissertation is on attention and visible properties. In this chapter, I defend the simple idea that we come to know about these properties by attending to them as they figure in visual experience.[54]

Understood in this way, attention is a means of fixing beliefs which inherit their content from visual experience: through attention, you come to believe that things have the very properties which meet the conditions of *Presence* in your visual experience. You believe that the cat has a certain shape and colour, say, because that shape and colour figure essentially in the phenomenology of your visual experience of the cat. Beliefs of this kind have a distinctive epistemic status. They inherit the distinctive reliability of visual processing (Dretske 1997), or the distinctive justification or warrant provided by conscious vision (McDowell 1994; Burge 2003).

As I said in Chapter 4, this epistemic role for attention is not obviously connected with effects on the determinacy of visual experience. Rather, it's naturally understood in terms of aspects of your overall experience over and above the distinctively visual or perceptual. Various proposals about the nature of non-perceptual attention are consistent with

---

[54] Recent work in both philosophy and cognitive science tends to focus on attention to physical objects or locations. But we can further specify what is attended in terms of the intrinsic properties instantiated by an object or at a location. Perhaps one could, in principle, attend to a property without attending to an object or location which instantiates it. However, I take it that the two forms of attention are not usually dissociated in our experience. See Huang and Pashler (2007) for an information-processing account of attention to properties that predicts this.

its playing this epistemic role. For example, we might understand this form of attention as a *sui generis* mode of consciousness, a way of selecting what you see in terms of which conscious thought and belief may be causally and rationally explained. Alternatively, we might propose that this form of attention itself consists in conscious cognition or thought. In Chapter 6 I will suggest some reasons for favouring the latter proposal. But in this chapter I focus instead on the epistemic role which attention plays – on the idea that attention is a significant, reliable source of visual knowledge, because it allows you to form beliefs which inherit their content from vision. I assess an empirical challenge to this epistemologically important idea.

The challenge lies in the fact that attention also plays a quite different role, in fixing the character and content of visual experience itself. James claimed that attention highlights attended properties by giving them a '*more intense*' appearance (1890: 425), and recent experiments confirm this. For example, cuing attention to a colour makes it appear more saturated; cuing attention to a gap in a shape makes it appear larger (Fuller and Carrasco 2005; Gobell and Carrasco 2006).

We should distinguish this from James's claim that attention makes visible properties *clearer*, which can be understood in terms of effects on the determinacy of the properties which visual experience represents (§4.1). Where a colour appears more saturated on one occasion than another, or a gap appears larger on one occasion than another, this is not just a matter of the colour or shape appearing more clearly. It is not just a matter of visual experience representing a more determinate shape or colour. Rather, on the face of it, the property must appear *different* on each occasion.

Yet as James remarked, '*the intensification … never seems to lead the judgement astray*' (426). The challenge is to understand how this could be. How could attention so alter visual experience without distorting the way visible properties show up in it, and making visual experience illusory? How could attention be a reliable source of visual knowledge, rather than a systematically misleading source of belief?

From the point of view of epistemology, those responsible for the recent experiments take a pessimistic view. Marisa Carrasco and her colleagues found that '*attention alters appearance*'; they conclude that attention produces '*nonveridical percepts … by emphasizing relevant details at the expense of a faithful representation*' (Carrasco et al. 2008: 1162). I will argue that this conclusion is not required, given that visual experience

represents determinable properties. There are transformations in property space, further to those which correspond to a property's becoming more determinate, in terms of which we can capture Carrasco's results without ascribing nonveridical experience. Attention may make gaps appear larger, colours appear more saturated, and so on, without making visual experience illusory. I will show how to interpret the experimental data in terms of veridical visual experiences, and I will argue that this interpretation is as empirically plausible as an account in which attention makes visual experience illusory. The empirical evidence is consistent with the idea that visual attention is a reliable source of knowledge.

There is also a broader issue here, about the relationship between experiment, judgement and visual experience. Carrasco's work is distinctive partly because she connects experiments in psychophysics with claims about the content of visual experience. The psychophysical data consist in subjects' discriminations between visible stimuli, where discrimination is a form of judgement. Psychophysicists use these data to investigate processes in the visual system. But it's unobvious how we should connect work in psychophysics with claims about the content of visual experience, or with our understanding of visual knowledge. I argue that, to do this in a principled way, we need to recognise the role of determinable properties in visual experience and visual discrimination.

5.2 Attention Alters Appearance: The Experiments

Carrasco's studies concern involuntary covert attention. Recall that, when attention is captured by a sudden movement or change in peripheral vision, there is a short-term effect on the visual system. In an effect which peaks around 100 milliseconds after attention is cued and disappears after another 20 ms or so, processing in the visual cortex occurs faster for stimuli at the location in question, and contrast-sensitivity is increased for stimuli there – i.e. visual-cortical responses become more sensitive to differences in those stimuli (Carrasco and McElree 2001; Treue 2000).

In a series of studies, Carrasco and her colleagues argue that this effect has a phenomenological upshot: visual experience is altered, highlighting properties at the cued location by exaggerating them; attention increases apparent brightness contrast, colour saturation, spatial frequency, gap size, motion coherence and flicker rate (Carrasco, Ling &

Read 2004; Fuller & Carrasco 2006; Liu, Fuller & Carrasco 2006; Gobell & Carrasco 2005; Montagna & Carrasco 2006).

I'll focus on colour saturation, gap size and spatial frequency. The size of a gap in a plane figure was measured in degrees of visual angle. If a gap is 5° wide, two imaginary straight lines reaching from its edges to the centre of the fovea form an angle of 5°. Recall that, for standard experimental stimuli, spatial frequency is measured in sinusoidal cycles of change per degree of visual angle (cpd). In Figure 19, spatial frequency increases from left to right. Intuitively, we can think of the spatial frequency of a stimulus as its rate of change across space. For example, where a stimulus has a higher spatial frequency, it typically has a more detailed pattern.

Figure 19

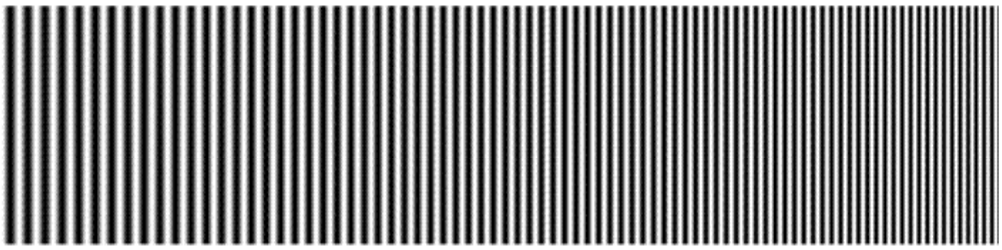

Figure 20 overleaf illustrates the key experimental paradigm. Subjects maintain fixation at a point in the centre of the display. A cue is flashed to capture their attention, either centrally or a little to one side. Soon afterwards, two stimuli appear for 40 milliseconds, a little to the right and left of fixation. Over many trials the saturation of the colour of one stimulus (the *test* stimulus) is varied, while that of the other stimulus (the *standard*) is held fixed. The task is to report the orientation of the stimulus which has the higher saturation. Carrasco uses similar experiments to investigate the effect of attention on the appearance of other properties, including gap size and spatial frequency.

The *Point of Subjective Equality* (PSE) is a value or level of a property, such as a gap size of 0.2° or a spatial frequency of 0.5 cpd. The PSE is defined as the value of the test stimulus such that, statistically, a subject is at chance in reporting the orientation of that stimulus: where the test stimulus has this value, the probability of the subject choosing it as

higher in the relevant value than the standard is 50%. For example, when attention is cued to the point of fixation, rather than to one of the stimuli, the PSE is roughly the value of the standard.

Figure 20

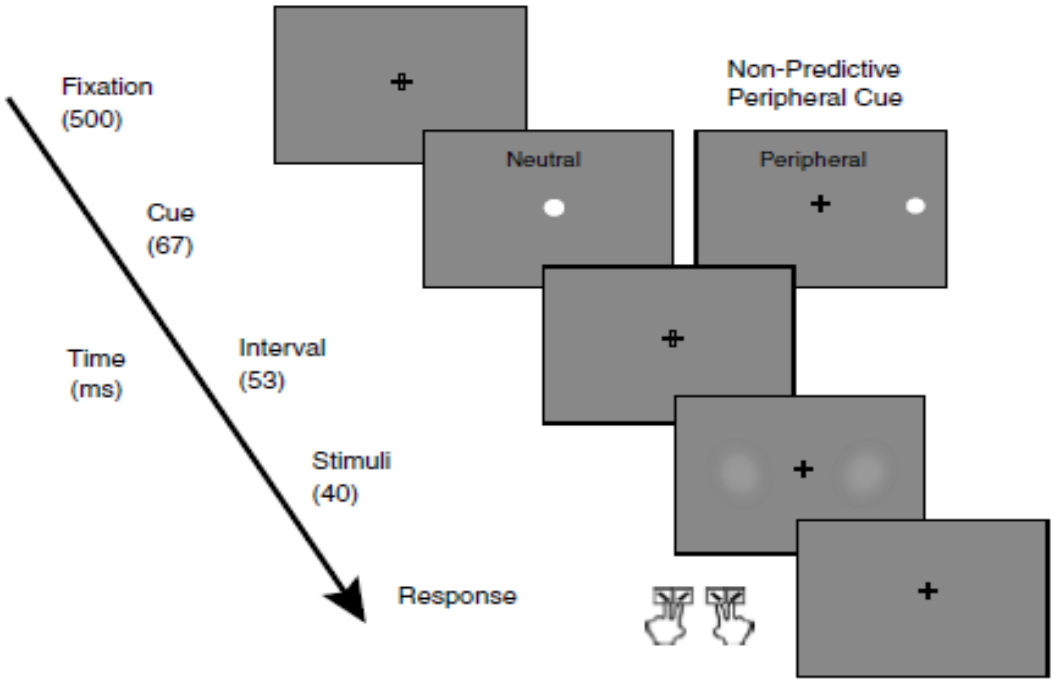

Reproduced from Fuller & Carrasco 2006: 4035, by permission of Marisa Carrasco & Elsevier (copyright 2006).

The crucial result is that the PSE was significantly lower when attention was cued to the test stimulus. For example, with the standard gap size at 0.23˚, the PSE was 0.20˚.[55] This is a significant, clearly visible difference in size. See Figure 21 to get a rough idea of the differences in spatial frequency (a) and gap size (b). Taking this together with the data about increased cortical sensitivity, Carrasco concludes that attention enhances or exaggerates the visual signal processed for an item, with the phenomenological effect that an attended item '*appears as if it were*' greater

[55] This is the PSE defined over a cohort of subjects, but similar results were found for individual subjects.

than it really is in gap size and so on (Carrasco et al. 2004: 311); attention produces '*non-veridical percepts*', rather than a '*faithful representation*' (Carrasco et al. 2008: 1162).

Figure 21

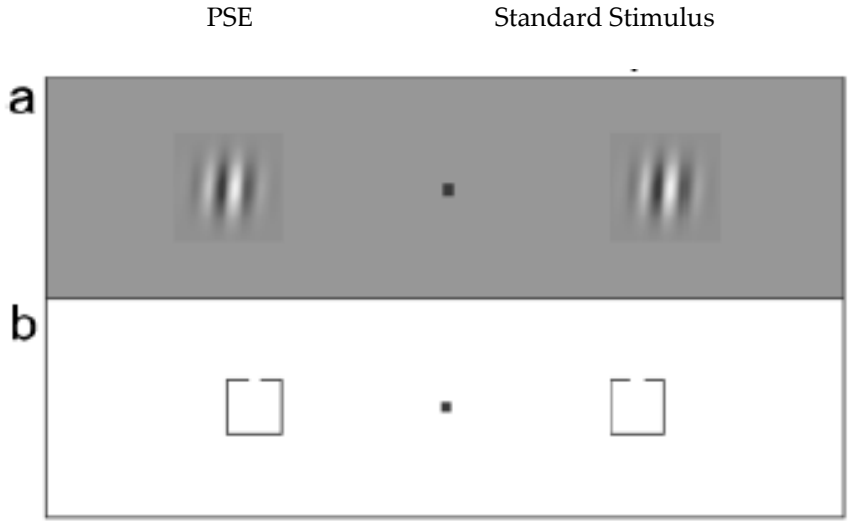

Reproduced from Gobell & Carrasco 2005: 650, by permission of Marisa Carrasco & the American Psychological Society (copyright 2005).

How does this threaten the epistemic role of attention? Here's an intuitive first pass at formulating the threat. In §5.1 I said that attention is a means of fixing beliefs which inherit their content from visual experience: through attention, you come to believe that things have the very properties which are represented in visual experience. Carrasco's experiments show that where you attend to an object, you experience it as having a property it doesn't really have – a higher saturation, a higher spatial frequency, a larger gap etc. So beliefs formed through visual attention turn out to be systematically false. They represent a colourful, detailed, gaping world, where by comparison reality is faded, bland and closed.

This first pass is not quite right. It conflates voluntary and involuntary attention. In Carrasco's experiments, subjects attend voluntarily to *both* stimuli, not just the stimulus that's cued. Subjects direct their thoughts to both stimuli, and perhaps they also select both of them with a pre-cognitive experiential focus, as discussed in §5.1. They have to

attend in this way to both stimuli, in order to compare them. Attention has the effect Carrasco describes only where it's cued exogenously. Voluntary attention also has pre-perceptual effects, but these differ systematically from the effects of involuntary attention (Prinzmetal and Landau 2008). For example, voluntary attention does not have the effect which Carrasco reports for involuntary attention, on either spatial properties or colours (Prinzmetal et al. 1998; Blaser et al. 1999).[56]

However, involuntary attention does affect the process of attending voluntarily which subjects exploit in their comparative judgements. Where attention is cued exogenously to one of the stimuli, this process is disrupted: visual experience is briefly distorted, and voluntary attention presents the subject with an exaggeration of the real properties before her.

So here's a more accurate way to formulate the threat. Where attention is cued exogenously, it *disrupts* the mechanism of experience and voluntary attention through which we form beliefs about visible properties. Visual experience then represents an exaggeration of the properties really present, and beliefs formed through voluntary attention do likewise. The disruption occurs when there is sudden change near a visible object. Since we inhabit a dynamic, changing environment, these circumstances often obtain. So attention routinely leads to false beliefs about visible properties. Much of what we took to be knowledge of visible properties turns out to be false belief.[57]

What's more, the threat may not be limited to cases in which involuntary attention actually distorts visual experience. In §5.1 I mentioned two general ways of understanding the idea that visual attention is a source of knowledge, in virtue of the fact that beliefs formed through attention inherit their content from visual experience:

[56] Liu et al. (2009) found that voluntary attention has an effect similar to that of involuntary attention on experience of one visible property: brightness contrast. If further work should show that voluntary attention in fact affects experience of spatial properties in a similar way, the problem I describe would differ, but my solution would still apply.

[57] Some philosophers argue that visual beliefs always abstract from the contents of visual experience. So we might hope that, even though involuntary attention changes visual representation, the changes are too minute to show up in false beliefs. However, subjects in the experiments *do* form false beliefs, as reflected in their false comparative judgements. Their beliefs don't abstract enough from the contents of experience to avoid being mistaken.

(1) Beliefs formed in this way inherit the distinctive reliability of visual processing.

Here the general threat is straightforward: the mechanism of attention and visual experience turns out to be an unreliable way to form beliefs about visible properties; visual attention is not a source of knowledge about visible properties.

(2) Beliefs formed in this way inherit the distinctive justification or warrant provided by visual experience and attention.

Here the status of Carrasco's threat is more controversial. We might argue that this justification or warrant is defeated, since visual experience and attention routinely ground false beliefs. I won't develop that argument here. Even if beliefs formed through visual attention *sometimes* constitute knowledge of visible properties, Carrasco's experiments threaten to show that in vastly many cases they do not.

Note that, so construed, Carrasco's threat to visual knowledge is not an instance of the traditional Argument from Illusion. Visual knowledge of properties is not here threatened by the mere existence of cases in which vision is illusory. Rather, the threat lies in the fact that the putative illusion is pervasive: according to Carrasco, involuntary attention *routinely* leads to illusion and consequent false belief.

5.3 Cue Bias and Mental Paint

Carrasco reports her data by saying the cued 0.20° gap '*appears as if it were*' 0.23° wide. Now there is a sense in which this is unobjectionable. The data consist in subjects' judgements, as expressed in their choices between the stimuli. When attention is cued to the 0.20° test gap, subjects judge it larger than the 0.23° standard gap in 50% of the experimental trials. This pattern of judgement mimics the pattern of judgement found when attention is not cued to either stimulus, and the test and standard are both 0.23° wide. On the basis of their visual experience, subjects judge in the same way whether they're faced with a cued 0.20° test gap, or with an uncued 0.23° test gap. In that sense, the cued 0.20° gap '*appears as if it were*' 0.23° wide.

However, it is a further question whether '*attention alters appearance*', where appearance consists in what's present or represented in visual experience. And it's a still further question whether attention alters the representational content of experience in a way that makes it '*nonveridical*'. We might seek alternative explanations of the pattern of judgement, explanations which don't involve illusory experience.

In this section, I briefly consider two such alternative explanations. I argue that neither of them successfully preserves the epistemic role of attention. Then in §5.4 I sketch my own alternative.

One natural idea here is that cuing attention might affect subjects' judgements directly, rather than via a change in visual experience (Prinzmetal et al. 2008). Carrasco claims that attention exaggerates the visual signal that's processed for the cued stimulus, with the effect that properties of the cued stimulus are exaggerated in visual experience. By contrast, on this alternative interpretation, cuing attention does not alter the visual signal or visual experience at all. Rather, just having their attention drawn to the stimulus makes subjects more likely to report on it. In the demanding experimental conditions, this is enough to explain the pattern of judgement.

Gobell and Carrasco (2005) provide convincing evidence that the pattern of judgement they found cannot be explained in this way. They conducted two control experiments. First, attention was cued to the location of the test stimulus after it had disappeared. If cuing influenced subjects' judgements directly, we could expect it to bias their judgements to the test even in this condition. But here the PSE did not shift as did in the original experiment.[58] Second, subjects had to report the stimulus which was *lower* in spatial frequency. Here the PSE shifted as it did in the original experiment: subjects were at chance in reporting on the cued test stimulus only where it had a significantly smaller gap. Cuing made subjects *less* likely to report on the test stimulus. So the result is not happily explained in terms of the idea that cuing attention directly biased subjects' judgements to the test.

[58] Prinzmetal et al. (2008) found the opposite result for post-cuing of a brightness-contrast stimulus: post-cuing shifted the PSE just as pre-cuing did. Carrasco et al. (2008) explain the discrepancy by showing that Prinzmetal et al.'s stimuli were near the threshold of invisibility, while their own were not. Carrasco et al. argue that cue bias *does* have a significant effect where stimuli are near-invisible, but that it does not have this effect for their stimuli.

What's more, cuing attention shifts the PSE only for some visible properties. For example, it shifts the PSE for colour saturation, but not for hue (Fuller and Carrasco 2006). It's hard to see how this could be explained by the idea that cuing attention has a direct effect on visual judgement. None of this is a demonstrative proof, of course, but the evidence does support Carrasco's claim that attention alters the visual signal and visual experience. I'll work with that claim from now on.

Ned Block (2010) takes a different approach to explaining the data without ascribing illusory visual experience. He claims that attention alters visual experience, but that attention does not alter the representational content of visual experience. He proposes that visual experience consists of a representational content together with a mode of presentation of that content. Modes of presentation play a distinctive phenomenological and psychological role in Block's account. The mode of presentation of an experience fixes its phenomenology, making the object of experience appear a certain way. For example, the mode of presentation of a subject's experience fixes a certain 'perceived gap size' for each stimulus – say a perceived gap size of 0.23° (Block 2010: 31). Nevertheless, the mode of presentation of an experience is not constituted by its representational content, and the mode of presentation may vary independently of the representational content. Where the mode of presentation fixes a perceived gap size of 0.23°, visual experience need not represent the gap as 0.23° wide. In that sense, visual experience is characterised by 'mental paint', a way things appear to a subject which is not a matter of how she represents them to be. What's more, this mental paint or mode of presentation may dictate the psychological role of a visual representation. In particular, Block explains Carrasco's data in terms of changes in the mode of presentation, brought about by cuing of attention.

For the Gobell and Carrasco experiment, Block's interpretation is as follows. Where a subject's attention is cued to a 0.20° gap, the mode of presentation in her experience of that gap is equivalent to the mode of presentation in her experience of an uncued 0.23° gap. Because the phenomenology of her experience of each gap is the same, she reports the cued 0.20° gap larger than the uncued 0.23° gap, in 50% of the experimental trials. On this approach, no effect of attention on the representational content of experience is required to explain the data. We can explain the data without ascribing systematic illusion to visual experience.

Block's proposal gives up on the idea that visual experience is a way of being open to properties of the environment, in the sense that you may come to know about them simply by believing that your environment instantiates them where they figure essentially in the phenomenology of visual experience. If Block's proposal is correct, then someone who simply believes that *P* is instantiated before her, where *P* figures essentially in the phenomenology of her visual experience, will not form beliefs in a reliable way. So if Block is correct, then visual experience and attention are not a distinctive source of knowledge in the way I have suggested.

Perhaps Block would not find this a fault with his account. His aim in discussing Carrasco's experiments is precisely to argue for the existence of 'mental paint'. However, the proposal also fails to explain how it could be that attention is not a systematically misleading source of belief. The changing mode of presentation, brought about by shifting attention, misleads subjects into making mistaken judgements. And if the misleading phenomenology of attentive experience leads to mistaken judgements here, won't it have the same effect elsewhere? Won't the exaggerated phenomenology of attentive experience systematically mislead us into believing that things have properties they don't really have?

Block appeals to James's discussion, to argue that the change in mode of representation need not in general lead to false beliefs, outside the specific experimental conditions:

> As we rightly perceive and name the same color under various lights, the same sound at various distances; so we seem to make an analogous sort of allowance for the varying amounts of attention with which objects are viewed; and whatever changes of feeling the attention may bring we charge, as it were, to the attention's account, and still perceive and conceive the object as the same.
>
> James 1890/1950: 426

James explains why altered visual experience does not lead to false belief in terms of a constancy effect: the perceiver grasps that the change is due to a shift in her perspective, a shift in her view on the world rather than in the world itself; so she does not attribute any change to the objects of experience. I don't deny that attention exhibits such a constancy effect. In

fact my own interpretation is consistent with this idea. But the challenge is to explain why this effect is not at work in the experiments, if it's at work in other cases. Subjects in the experiments *are* misled.

Block comments: 'Many experiments on perception in effect encourage the subject to judge appearance', where 'appearance' refers to the phenomenological mode of presentation, rather than the representational content of experience (54). But if anything Carrasco's experiments discourage judgements about the character of experience, as opposed to judgements about the objects of experience. Subjects were instructed to 'report the orientation of the Gabor of higher spatial frequency' (Gobell and Carrasco 2005: 645).[59] The instruction focuses on the actual properties of the stimulus, rather than its subjective appearance. To complete their task successfully, subjects must keep track of the objective spatial frequency of the stimulus. Furthermore, since the instruction focuses most immediately on orientation rather than spatial frequency, there's no reason for subjects to assume, confusedly, that what the experimenters are really asking about is the character of their experience of spatial frequency. So Block's proposal about mental paint does not explain how it can be that, in general, attention '*never seems to lead the judgement astray*' (James 1890/1950: 426), while attention does lead judgement astray in the experiments.

The interpretation I propose does explain this, and it does so in terms of changes in the representational content of visual experience. On my interpretation, we can explain the data consistently with the claim that visual attention is a reliable, distinctive source of knowledge in the way I explained above: through attention, you come to know that something has the property *P*, in virtue of the fact that *P* figures essentially in the phenomenology of your visual experience.[60]

[59] The type of stimulus used in known as a Gabor patch.

[60] My position here is consistent with Block's claim that there is mental paint. However, my position is inconsistent with his (2010) argument for that claim. Block makes a subtle empirical and philosophical case for denying that attention leads to illusions in Carrasco's experiments. He then argues, by eliminating alternative interpretations, that the data must therefore be explained in terms of mental paint. I defend an alternative interpretation.

5.4 Attention Alters Determinable Appearances

Let's take stock. I propose that we accept the following three claims made by Carrasco (for the example of gap size):

(1) Cuing attention to a 0.20° gap changes the visual signal processed for it, replicating the signal for an uncued 0.23° gap.
(2) As a result, visual experience of the cued 0.20° gap replicates visual experience of the uncued 0.23° gap. In this sense, cuing attention to the 0.20° gap gives it the same appearance as an uncued 0.23° gap.
(3) This effect of attention on visual experience is an effect on its representational content: visual experience represents the cued 0.20° gap in the same way as it represents the uncued 0.23° gap.

However, it doesn't follow that:

(4) Visual experience of the cued 0.20° gap is illusory.

On some assignments of representational content to visual experience, clearly (4) would follow. Suppose that visual experience represents exact gap sizes. It represents the uncued 0.23° gap as exactly 0.23° wide. Then if cuing attention to the 0.20° gap replicates this experience, cuing attention makes visual experience illusory.

However, visual experience does not represent exact or maximally determinate gap sizes. More generally, visual experience often represents determinable properties without representing determinations of them. Now on the analysis I explained in Chapter 2, the determinables determined by one more determinate property may vary in at least two ways, which I'll call *dilation* and *translation*, corresponding to two kinds of transformation in property space.

Where properties *A* and *B* vary by dilation, we can model the difference between them as a dilation in property space: to map the region of property space corresponding to *A* onto the region corresponding to *B*, we alter the size of the region without altering its shape. So one of these properties is more determinate than the other. For example, consider the very simple example of width, which has a one-dimensional property space: the properties *between 0.19° and 0.27° wide* and *between 0.21° and 0.25° wide* are dilations of one another. Properties with richer property spaces, for example shapes and colours, may equally vary in this way.

Dilation is the change in represented properties that's most naturally associated with attention, and as we've seen involuntary attention does induce dilation in the spatial properties which visual experience represents. Indeed, the experiments at issue in this chapter do plausibly involve dilation with respect to some of the properties which visual experience represents. Fuller and Carrasco suggest that 'increased contrast and saturation facilitate the discrimination of features of the signal, and make it easier to discriminate the signal from the background' (2006: 4043). For example the shifting PSE for saturation, produced by attentional cuing, was correlated with steadily increasing accuracy in subjects' comparative judgements about hue (4042).[61] This suggests a gradual effect on the appearance of hue, consistent with the idea that attention here leads to a dilation in represented hue: a signal representing exaggerated saturation allows your visual system to compute a more determinate hue – a more determinate shade of scarlet, say, rather than just scarlet or even red.

But on Carrasco's interpretation, this advantage comes at a cost to the veridicality of experience: you see the thing's hue more determinately, at the cost of seeing a saturation which is not in fact present. In general terms, the trouble is that evolutionary advantage and truth are fare-weather friends. It might be advantageous to see the shape or hue of a suddenly-moving object more determinately, even at the cost of a false lemma about the thing's contrast or saturation. And we cannot explain Carrasco's data about increased apparent hue, or increased apparent gap-size and so on, in terms of dilation in the properties which visual experience represents. For a gap to appear larger is not for its size to be represented more determinately; for a colour to appear more saturated is not for it to be represented more determinately.

However, determinable properties may also vary by translation. Where properties *A* and *B* vary in this way, we can model the difference between them as a translation in property space: to map the region of property space corresponding to *A* onto the region corresponding to *B*, we shift the region through property space without altering its size. So the properties differ, but neither is more determinate than the other. Where

[61] The PSE for hue itself did not shift with attentional cuing. Fuller and Carrasco speculate that there is evolutionary advantage in shifting PSEs only for properties which admit of quantitative, rather than qualitative, variation: exaggerated saturation in a hue facilitates discrimination of the hue; distorted hue would have no parallel advantage.

two properties are translations of one another, they may nonetheless share some determinations: there may be one property which determines each of them. For example, the widths *between 0.16° and 0.24°* and *between 0.19° and 0.27° wide* vary by translation.

Suppose that, when attention is not cued to a gap of 0.20°, a subject's experience represents it as between 0.16° and 0.24° wide. And suppose that, when attention is not cued to a gap of 0.23°, her experience represents it as between 0.19° and 0.27° wide. If cuing attention to a 0.20° gap replicates this experience, as Carrasco claims, the cued 0.20° gap will likewise be represented as between 0.19° and 0.27° wide: cuing attention will induce translation in the property which visual experience represents. Where the two gaps are represented in this way, it's compatible with the content of experience that either gap be larger. But Carrasco's task is *forced choice*: subjects have to report one of the gaps as larger, even if they're not confident in their judgement. So where the cued test gap and the uncued standard gap are each represented as between 0.19° and 0.27°, we should expect the pattern of judgement which Carrasco found: subjects report the test gap larger around 50% of the time. And on this interpretation, visual experience of the cued gap is not illusory, because although the properties represented in attentive and inattentive conditions translate, they are each determined by the actual determinate width of the gap. A gap of 0.20° really is between 0.19° and 0.27° wide.

Similar interpretations can be applied with respect to the other visible properties Carrasco and her colleagues studied, including spatial frequency and colour. We can explain the data about attention and visible properties using an interpretation of this form. And we can also understand James's observation in this way, at least as it applies to involuntary attention. James said that attention gives visible properties a '*more intense*' appearance. On the interpretation of the data I'm proposing, we can understand this as follows. Cuing attention to a property gives it an appearance which an uncued, more intense property would have. Your visual experience represents the cued property in a way in which it would also represent an uncued, more intense property. Still, the altered appearance need not be illusory. What's altered is the determinable property which visual experience represents, and different veridical experiences of a property may represent its different determinables. For example, uncued experience of the 0.20° gap might represent it as between 0.16° and 0.24° wide, while cued experience represents it as between 0.19° and 0.27° wide. Both experiences are veridical.

This interpretation explains why subjects are misled in Carrasco's experiments, even though, in general, attention '*never seems to lead the judgement astray*'. Subjects are forced to make a judgement which is underdetermined by their experience. For example, the content of their experience is compatible with either gap being larger. In light of the representational content of their experiences, it's equally probable that either gap be larger. This feature of their experiences is what misleads them, but it's misleading only with respect to the specific, demanding task which Carrasco's subjects are given. By contrast, where a subject's beliefs about the gaps inherit their content from her visual experience, her beliefs are true. This interpretation is consistent with the claim that visual attention is a reliable source of knowledge.

Block comments that 'the change invoked by changing attention does not look like a change in the world' (2010: 53). This is surely right. When your attention is cued to a gap, the gap does not seem to grow. My interpretation predicts this. If you represent something as having two different but compatible determinable properties, over the course of a visual episode, this is not ordinarily grounds for supposing that the object has changed.

I don't suppose that the particular assignment of representational content I've used is the correct one. For one thing, with respect to spatial properties attention induces both translation and dilation. But the interpretation works for various assignments of content. The basic constraint is that experience must be sufficiently indeterminate: the determinables represented must correspond to a sufficiently broad range of determinate gap sizes. For example, if the uncued 0.20° gap were represented as between 0.19° and 0.21°, cuing attention to it could not replicate the appearance of the 0.23° gap in the way I've described, and remain veridical.[62]

In Carrasco's experiments, it is plausible that the determinable properties represented correspond to fairly broad ranges of determinates, because it's plausible that visual resolution may be quite limited. Because the experiments are designed to probe the short-term effects of involuntary attention, stimuli are presented extremely briefly – for only

[62] There is also a further constraint: the actual determinate size of the gap must not be near the bottom of the range corresponding to the determinable represented in uncued experience. For example, if the uncued 0.20° gap were represented as between 0.19° and 0.27° wide, intensifying visual experience in the way I've described would result in an illusion. I discuss this below.

40ms. Neurons in the visual cortex become sensitive to increasingly higher spatial frequencies over the course of a normal fixation of around 200ms (Frazor et al. 2004). 42ms after onset of a stimulus, the maximum spatial frequency to which cells in V1 are sensitive is 2½ times lower than that to which they're sensitive 74ms after onset (Bredfeldt and Ringach 2002). And vision scientists connect these facts about early vision with conscious vision, arguing for a distinction between vision 'at a glance' and 'vision with scrutiny' (Hochstein and Ahissar 2002). On this approach, conscious vision is at first characterized by the schematic, categorizing information encoded in higher visual areas; only later, once feedback from higher areas tunes processing in the lower areas, do details of the scene reach consciousness.

It is also plausible from the point of view of phenomenology that subjects in Carrasco's experiments experience determinable properties corresponding to quite broad ranges of determinates. For example, Maurice Merleau-Ponty remarked on the indeterminacy of experience early in a visual episode, commenting that normal-length visual episodes are 'a passage from the indeterminate to the determinate' (1945/1962: 36). And we need not take his word for it. Try Carrasco's task for yourself: http://philosophy.berkeley.edu/people/page/99. Even at 100 ms, 2½ times longer than Carrasco's subjects get, it's hard to make out determinate detail. In principle, one might insist that the short temporal window limits only the determinacy of judgement, rather than the determinacy of visual experience. But that proposal is not positively supported either by the phenomenology of the task or by the neuroscientific evidence. It is consistent with both that subjects in the experiments experience determinable properties corresponding to broad ranges of determinates. None of this *demonstrates* that subjects in the experiments experience determinables as broad as my interpretation requires. But it does make that an empirically relevant possibility, provided we can explain subjects' powers of discrimination in terms of such experiences.[63]

[63] Block allows that visual experience might represent determinables, but he gives two reasons for thinking that the data cannot be explained in these terms: (i) representations of determinables can't capture the phenomenology of experience, the conscious awareness to which subjects in the experiments respond; (ii) these representations cannot explain the accurate fine-grained discriminations subjects make when their attention is not manipulated. I addressed the first point in Chapter 2. I address the second next.

5.5 Discriminating the Determinable

Carrasco's data consist in her subjects' attempted discriminations between stimuli. Subjects have to discriminate the stimuli according to an ordering, with respect to colour saturation, gap size, spatial frequency etc. That is, they have to report which of the stimuli is more saturated in colour, which contains the wider gap, which is higher in frequency, and so on. They discriminate fairly accurately when their attention is not cued to either stimulus. For example, in the experiments about gap size and spatial frequency, they discriminated accurately on 9/10 trials overall (Gobell and Carrasco 2005). To see how this is consistent with an interpretation along the lines I have proposed, we need to reflect a little on the relationship between visual discrimination and visual experience.

There is a sense in which discrimination is the most basic visual capacity: in order to detect the various aspects of a visible scene, the visual system must respond differentially to different objects and to their different features. I want to set aside that sense of 'discrimination', because it's not the sense in which Carrasco's data consist in attempted discriminations. Carrasco's subjects make judgements as to which stimulus is more saturated, which has a larger gap, and so on. The data consist in these judgements, as expressed in the pressing of a button. So in the sense that's relevant here, discrimination is a form of judgement or knowledge, an achievement of a person rather than a sub-personal achievement of the visual system. To discriminate *a* from *b*, in the relevant sense, is to recognise that *a* and *b* are distinct (Williamson 1990). Visually to discriminate *a* from *b* is to recognise that *a* and *b* are distinct on the basis of visual experience. We discriminate things with respect to their properties. For example, you might discriminate *a* from *b* by recognising that *a* has a length which *b* does not have. We can also talk of discriminating the properties of things – for example recognising that the length of *a* and the length of *b* are distinct.

To understand visual discrimination, we need to understand that visual experience of a property *P* and visual discrimination with respect to *P* are distinct, contingently related achievements. In particular:

(i) The properties which visual experience represents *a* and *b* as having may differ, even where the subject cannot discriminate *a* from *b* with respect to those properties.

&

(ii) A subject may discriminate *a* from *b* with respect to some properties, even where the distinctness of these properties is not deducible from the contents of her experiences.

I'll explain (i) first, and in doing so I'll defend an account of visual discrimination which accommodates visual indeterminacy. Then I'll apply this account in explaining (ii), which bears directly on Carrasco's experiments.

(i) Determinable properties have a useful feature: within a level, so to speak, determinations of a property are incompatible with one another. More precisely:

Exclusion

> If *Q* and *R* each determine *P*, and there is no *S* such that *S* determines both *Q* and *R*, then nothing could have both *Q* and *R* simultaneously, in the same location, and in the same respect.

Where *Q* and *R* meet this condition, I will say that *Q* and *R* are *incompatible*. In a property space, this is reflected by the fact that *Q* and *R* correspond to non-overlapping regions. For example, nothing can be coloured in such a way that it is both blue and red all over at the same time, and in the same respect; nothing could be both between 9.995 and 10.005 cm long and between 10.008 and 10.018 cm long simultaneously, in the same location, and in the same respect.

We exploit *Exclusion* to discriminate things with respect to their determinable properties. You can discriminate *a* from *b*, with respect to a determinable property *P*, if you recognise that *a* and *b* have incompatible determinations of *P*. In the simplest cases, we do this according to the following rule for visual discrimination:

The Canonical Method

> Judge *a* distinct from *b* with respect to *P*, where your visual experience represents *a* and *b* as having incompatible determinations of *P*.

For example, suppose you're presented with two clearly visible lines. Your visual experience represents line *a* as between 9.995 and 10.005 cm long, and it represents line *b* as between 10.008 and 10.018 cm long. In this situation, and all else being equal, you can visually discriminate *a* from *b* with respect to their lengths: the lengths your experience represents *a* and *b* respectively as having are incompatible with one another. So you can recognise that *a* is distinct from *b*: whatever determinate lengths they have, these lengths are not identical; there is a length which *a* has and *b* lacks (and *vice versa*).

Therefore, differences in the properties which visual experience represents are one thing, discrimination with respect to those properties another. The two should not be conflated. This is clearest in cases where your visual experience represents *a* and *b* as having different but compatible determinations of *P*: here you may not be able visually to discriminate *a* from *b* with respect to *P*. For example, where your visual experience represents *a* and *b* as having two determinable lengths which share some determinations, you may not be able to discriminate *a* from *b* with respect to their lengths. If your experience represents *a* as between 9.995 and 10.005 cm long, and it represents *b* as between 10.003 and 10.013 cm long, the lengths your experience represents *a* and *b* respectively as having are consistent with one another; for all the contents of your experience entail, *a* and *b* might share a length.[64]

I take it that some philosophers will object to this picture of visual discrimination, even once the indeterminacy of visual experience is granted. For example, Christopher Peacocke (1989) allows that visual experience represents inexact directions, but still assumes the following: if visual experience represents precise (though inexact) directions, then the directions which a subject's visual experience represents *a* and *b* as having must be identical, if the directions of *a* and *b* are indiscriminable to her.

Peacocke explicitly motivates this assumption only by appeal to a principle that 'directions experienced … seem to be the same if and only if they are the same' (301). This is rather ambiguous. First, it's not clear in what sense indiscriminable directions must seem to be the same. That doesn't obviously follow from their not recognisably differing. Second, it's not clear how to connect the visual indiscriminability of two directions with Peacocke's idea that the 'directions experienced' – the directions

[64] Cf. Hellie's (2005) account of indiscriminability for the colours. Hellie exploits the notion of overlap in colour space, but not the notion of determination.

which experience represents – seem to be the same. On the face of it, for visible directions to seem the same or different is one thing, while for visual *experiences* of direction to seem the same or different is another.[65]

However, we might reconstruct the thinking here in terms of two plausible-sounding principles:

Transparency

> A subject knows through introspection that her conscious visual representations have a property *P*, only if the objects of those representations seem to her to have *P*.

Access

> A subject's conscious visual representations differ only if she is in a position to know through introspection that they differ.

One property which both representations and objects may have is that of qualitative difference. So together, *Transparency* and *Access* guarantee that where a subject's visual experiences represent different properties, she is in a position to judge, on the basis of visual experience, that the objects of her experience differ: contrary to the account I have sketched, any difference in the properties which visual experience represents guarantees the appearance of a difference between the objects of experience.

Transparency expresses the thought that everything which introspection reveals about visual representation, visual representation reveals about the visible world. It's a generalization of G.E. Moore's famous comment that introspecting the experience of blue reveals only the experienced blue. You could motivate *Access* without holding that we

[65] The expression 'direction experienced' is ambiguous between a direction which is an object of experience, and a direction which visual experience represents something as having. For example, a direction of 50° might be the object of an experience which represents something as having the direction *between 45° and 55°*. Peacocke is concerned with 'direction experienced' in the second sense – with 'the contents of the experiences' (301). Read in the first sense, his principle would be falsified by any illusory experience of direction. But he trades on the ambiguity when he connects his principle with visual indiscriminability: when two directions are visually indiscriminable, he takes it that visual representations of those directions are likewise indiscriminable. In what follows, I make this step explicit in terms of *Transparency*.

know all about our experiences. Instead, you might motivate it on somewhat Fregean grounds: where conscious visual representations differ, this difference must be cognitively significant; the difference will be cognitively significant for someone who is able to recognise it.[66]

However, these motivations are far from decisive. It's not at all obvious that Moore's observation should be generalised into *Transparency*. Perhaps consciousness of first-order visible properties such as colours and shapes is introspectible only where some objects of experience seem to have those properties. But we can accept this without also accepting a parallel principle for second-order properties, such as the property of differing. Certainly Moore's observation does not by itself motivate the parallel principle.

Similarly, *Access* is hardly required by the demand that differences in visual representation be cognitively significant. Differences in visual representation might be cognitively significant for someone unable to recognise them. On the face of it, differences in visual representation might be cognitively significant for someone altogether unaware of her own experiences. What's required is just that she respond differently to the *objects* of experience, in virtue of the different properties which her visual experience represents. And we have principled reason to ascribe differences in visual representation independently of their impact on visual discrimination. The evidence I explained in Chapter 2 includes facts about the sensitivity of the visual system which we can probe independently of visual discriminations. For example, some of the evidence about spatial resolution in early vision comes not from discrimination behaviour on the part of human subjects, but from single-cell recordings in Macaque monkeys – whose early visual system is substantially like ours (DeValois and DeValois 1988). Presumably, differences in visual representation are cognitively significant in virtue of the downstream effects of visual experience, rather than the earlier processing which makes it possible. But here too the evidence outstrips the facts about visual discrimination. For example, the length which your visual experience represents an object as having may explain how you go to grasp the object; the apparent saturation of a colour may affect capacity

[66] This may be Peacocke's motivation. Later in the paper (1989: 308), he draws an analogy between his principle about experienced directions and a parallel principle about the directions represented in Fregean senses. We might accept the latter without accepting the former, but that is not by concern here.

to identify its hue (Fuller and Carrasco 2006), or your evaluative response to the colour (Palmer and Schloss 2010).

Moreover, there are compelling arguments against the conjunction of *Transparency* and *Access*. First, because visual experience is indeterminate, if *Transparency* and *Access* were true, then visual discriminations could not in general be justified by visual experiences. Where visual experience represents *a* as having a determinable property *P* and *b* as having a different but compatible determinable property *Q*, someone simply taking her experience at face value is not justified in taking *a* and *b* to differ with respect to these properties. For all her experience reveals, *a* and *b* might share a determinate property which falls within the intersection of the property space for *P* and the property space for *Q*. For example, if your experience represents *a* as between 9.995 and 10.005 cm long, and it represents *b* as between 10.003 and 10.013 cm long, the lengths your experience represents *a* and *b* respectively as having are consistent with one another. But together, *Transparency* and *Access* require you to judge *a* and *b* distinct in length.

Second, *Transparency* and *Access* lead quickly to contradiction, when combined with some manifest facts about visual discrimination. As Peacocke notes, there are cases in which – with respect to colour, length, and other spatial properties – *a* is indiscriminable from *b*, and *b* is indiscriminable from *c*, yet *c* is discriminable from *a*. (If you find this less than manifest, add as many items as you like to the series. What's important is just that only two of them be discriminable.) Now the conjunction of *Transparency* and *Access* guarantees that two objects are visually indiscriminable to you with respect to length (say), if and only if your visual experience represents each of them as having some length *L*. So together *Transparency* and *Access* entail that your visual experience takes the following impossible form: *a* is represented as *L*, *b* is represented as *L*, *c* is represented as *L*, but *a* and *c* are not represented as the same length.[67]

[67] Peacocke himself makes effectively this point. But to avoid the contradiction, he claims that visual experience does not represent precise colours or spatial properties (even precise though inexact colours or spatial properties). On this basis, he introduces a more complex way of understanding the contents of experience, such that no unique colour or spatial property meets the conditions of *Presence* for a subject. Given the theoretical value of *Presence*, and given that *Transparency* and *Access* are each independently suspect, I prefer to reject Peacocke's principle about visual discrimination.

I will take it, then, that visual discrimination standardly operates according to the *Canonical Method*. This is not to say that we are able to discriminate objects wherever visual experience represents them as having incompatible properties. There may be cases in which visual experience represents *a* and *b* as having incompatible properties, but further conditions on knowledge fail. For example, consider this case:

Borderline

> *a* has the length *10.0049 cm*.
> *b* has the length *10.0051 cm*.
> Visual experience *E* is sensitive to the length *between 9.995 and 10.005 cm*, and *E* represents *a* as having that length.
> Visual experience *F* is sensitive to the length *between 10.005 and 10.015 cm*, and *F* represents *b* as having that length.

For the purposes of this example, read those ranges of lengths non-inclusively – so *E* and *F* represent *a* and *b* as having incompatible lengths. It has been suggested to me that cases of this kind pose a problem for my account of the relationship between visual experience and visual discrimination. Surely *a* and *b* are not discriminable with respect to those lengths. Doesn't my account cast them as discriminable? Is the account not therefore inconsistent with the manifest facts about visual discriminability?[68]

The account does not predict that *a* and *b* are discriminable with respect to length in *Borderline*. The circumstances in *Borderline*, though normal, are quite extreme. On the one hand, *a* has a determinate length at the upper extreme of the lengths to which *E* is sensitive. So circumstances are at the upper extreme of circumstances in which *E* occurs in the normal way. On the other hand, *b* has a length at the lower extreme of the lengths to which *F* is sensitive. So circumstances are at the lower extreme of circumstances in which *F* occurs in the normal way. Perhaps this combination of circumstances is altogether impossible, for visual experiences of length. It's not clear whether any normal circumstances could lead to such divergent effects on visual responses to *a* and *b*. But let that pass: perhaps *Borderline* is in principle possible, in a case in which your viewing angle or the lighting, say, is very unusual.

[68] Particular thanks to Michael Caie and John Campbell for pressing me here.

Still, under circumstances so extreme, where visual experience represents *a* and *b* as having incompatible lengths, this does not suffice for knowledge that *a* and *b* have different lengths. Had circumstances been just a little different, you would have been subject to visual illusion: circumstances would have been abnormal in such a way that you enjoyed experience *E*, say, while the length to which *E* is sensitive was not instantiated. So if you judge *a* and *b* distinct in length, on the basis of the evidence in *Borderline*, you judge in a way which could easily have gone wrong. You fail standard conditions on knowledge. If you had judged in this way in slightly different circumstances, you would have judged *a* and *b* distinct in length even though they were identical in length – so you fail the safety condition of Williamson 2002. And if *a* and *b* had not differed in length, you would have judged them distinct all the same – so you fail the tracking condition of Nozick 1981.

I take it that, where visual experience represents *a* and *b* as having incompatible lengths, *a* and *b* look different with respect to length. So *Borderline* is a case in which *a* and *b* look different with respect to length. Furthermore, it is a case in which, if you take your experiences at face value, you will judge truly that *a* and *b* have distinct lengths. However, the circumstances are sufficiently extreme, and so sufficiently risky, that you should withhold judgement as to whether *a* and *b* have distinct lengths. At any rate, if you do go ahead and judge them to have distinct lengths, the risk is such that your true judgement does not amount to knowledge. In *Borderline*, *a* and *b* are not discriminable with respect to length. On the other hand, where circumstances are not so risky, the *Canonical Method* is a reliable way to discriminate colours and spatial properties.

(ii) On the interpretation I propose, the *Canonical Method* will not be available to Carrasco's subjects, in trials where the stimuli are similar. For example, if a subject represents the 0.20° gap as between 0.16° and 0.24° wide, and the 0.23° gap as between 0.19° and 0.27° wide, she cannot discriminate them by the canonical method. Yet subjects discriminate fairly accurately when their attention is not cued to either stimulus. Overall, they discriminate accurately on 9/10 trials. How do they do it?

Psychophysicists don't usually interpret their data in terms of visual experience or its representational content. Typically they use *signal detection theory*. Signal detection theory models visual judgement as the output of a decision process, the input to which is a response in the visual

system. Visual processes encode the (subjective) probability that the stimulus with which you're presented has a certain character, conditional on a certain visual response. The process by which you reach a judgement or decision about the character of the stimulus is modelled as a probabilistic inference, based on the visual response. In Carrasco's experiment, the gap in each stimulus produces a visual response. Depending on the relative levels of these responses, the visual system encodes a probability that the test gap is larger, and a probability that the standard gap is larger. On this basis, the system computes a decision as to which gap is larger.[69]

Signal detection theory makes no mention of visual experience, but it's open to us to model visual experience as a visual response. From this perspective, we can see that subjects need not be able to discriminate stimuli by the canonical method, in order to discriminate them accurately. Take a subject who experiences the 0.20° test gap as between 0.16° and 0.24° wide, and the 0.23° standard gap as between 0.19° and 0.27° wide. If she's to discriminate accurately, it must be probable that the standard gap is larger, conditional on these responses, and she must be sensitive to this fact.

Conditional on these responses, it *is* probable that the standard gap is larger. Recall that, where a visual representation *R* represents a property *P*, *R* is reliably correlated with *P* in circumstances such that *R* occurs in the normal way. For example, where visual experience represents the test gap as between 0.16° and 0.24°, there are different possible circumstances in which this experience occurs in the normal way, such that the stimulus has any gap size within this range. And where experience represents the standard gap as between 0.19° and 0.27°, there are different possible circumstances in which this experience occurs in the normal way, such that the stimulus has any gap size within this range. Taking all these possible circumstances together, there are more possibilities such that the standard gap is wider than there are possibilities such that the test gap is wider.[70]

Carrasco's subjects are trained in the task. Before results are recorded, they go through large blocks of trials, getting feedback on each

[69] See MacMillan and Creelman 2005: 113-4, for a model of experiments of this form.

[70] Strictly speaking, the probability distributions are probably Gaussian, and so weighted to widths near the centre of the range. But it's fair to assume that this weighting is similar for the test and the standard. So it's still more probable that the standard gap is wider, than it is that the test gap is wider.

trial as to whether they judged correctly. And their performance improves radically after training. Training gives subjects the opportunity to become sensitive to the probability that a particular gap is wider, conditional on their visual experiences of the gaps. For example, where you represent one gap as between 0.16° and 0.24° and the other as between 0.19° and 0.27°, you'll judge correctly most of the time if you say the latter gap is wider; feedback about your judgements informs you of this fact. We should not think of subjects as explicitly or consciously learning about the relevant probability-distributions. In general, psychophysicists don't take us to be aware of the subjective probabilities on which our perceptual judgements depend. Rather, these probabilities are encoded implicitly in sub-personal systems. From the subject's point of view, each gap has a certain visual appearance. On the basis of these appearances, she estimates which gap is larger, even where she's not at all confident in her judgement. Signal detection theory gives us a way of understanding how this could be a reliable method of discrimination.[71]

Consider again the case in which subjects' judgements are systematically *in*accurate: where their attention is cued to the test stimulus. Carrasco claims that subjects experience the cued 0.20° gap in the same way as they experience the uncued 0.23° gap. For example, a subject's visual experience might represent both as between 0.19° and 0.27° wide. Conditional on these responses, it is equally probable that either stimulus is wider. If she judges in the way learned during training, the subject will judge the 0.20° gap wider in roughly 50% of trials. That explains the shifting PSEs brought about by cuing attention. Yet the subject's experience is veridical.

This is of course just a rough sketch of how we might explain the accuracies and inaccuracies of Carrasco's subjects' judgements. Certainly I haven't offered a statistical argument for any specific reading of the data. But the interpretation I have proposed works for various assignments of representational content. For simplicity's sake, I used an example in which the determinate gap size of stimulus (0.23°) is at the centre of the range of gap sizes consistent with the visual representation (0.19°-0.27°). The interpretation does not turn on this. Realistically, subjects will represent

[71] Just anecdotally, subjects in similar experiments at Michael Silver's Berkeley lab said they were often guessing, even when they judged correctly. That was my own experience too. One researcher told me that, when he first programmed the experiment, he thought he'd done it wrong because the task seemed impossible. He hadn't, and trained subjects 'guessed' very accurately.

different determinables in different cases, with the actual property of the stimulus at various points within the range. In cases in which this point is near the edge of the range consistent with their visual representation, subjects are liable to judge inaccurately, even when attention is not cued. For example, if you represent the 0.20° test gap as between 0.19° and 0.27° wide, and the 0.23° standard likewise as between 0.19° and 0.27°, you're liable to get it wrong about half the time. But statistically these cases are unlikely to be the norm.

Similarly, where the determinate gap size of the stimulus is near the bottom of the range consistent with uncued visual representation, cuing attention is liable to produce illusions. For example, if uncued experience represents the 0.20° test gap as between 0.19° and 0.27° wide, intensifying visual experience in the way I've described will produce an illusion. But again, statistically these cases are unlikely to be the norm. Even if attention occasionally causes visual illusions, we need not suppose that it does so routinely.

Signal detection theory is relevant to these questions about veridical experience, only because visual experience represents determinable properties. For example, suppose that visual experience represented maximally determinate gap sizes. If that were right, then only a misrepresentation could make it subjectively probable that the cued 0.20° gap was wider than an uncued 0.23° gap.

In general, to understand what psychophysical experiments tell us about the content of visual experience, we need to recognise the role of determinable properties in visual experience and visual discrimination. Psychophysics reveals that various contextual factors other than the character of a stimulus affect the course of visual processing. In demanding tasks, these effects sometimes lead to systematically inaccurate judgements (see Carrasco et al. 2008). If we assume that visual experience represents maximally determinate properties, the experiments seem to reveal systematic illusion. The idea that visual experience makes us simply open to our environment looks like a naïve mistake. In fact, however, different veridical experiences may represent one unchanging property differently, where they represent different determinables of that property. As a result, contextual effects may alter the content of visual experience, leading to inaccurate judgements in demanding tasks, without experience being illusory. Where beliefs simply inherit their content from visual experience, experience may be a reliable source of knowledge. In

that sense, it's consistent with the findings of psychophysics that visual experience is a way of being simply open to our environment.

In particular, because visual experience represents determinable properties, attention may change the way properties appear in visual experience, without generating illusions. James claimed that attention gives visible properties a 'more intense' appearance. Where visual experience represents determinable properties, we can understand this in a way that's consistent with the view that attention plays a distinctive epistemic role, as a means of fixing beliefs which inherit a distinctive reliability or justification from visual experience. Attention to a property gives it an appearance which an unattended, more intense property would also have. Nonetheless, the attentive appearance need not be illusory. For all the data show, attention may well be a reliable source of visual knowledge.

Chapter 6

ATTENTION AND THE LIMITS OF OPENNESS

6.1 Openness

Visual experience is the distinctively visual aspect of an overall experience in virtue of which you see (§1.1). In Chapter 5, I defended the claim that this aspect of experience makes you open to the visible environment in the following sense:

Epistemic Openness

> A subject may come to know that her environment instantiates a property *P*, by believing that her environment instantiates *P* where *P* figures essentially in the phenomenology of her normally-occurring visual experience.

In virtue of this form of openness, visual experience is a distinctive source of knowledge.

Philosophers have also been interested in the proposal that visual experience is a way of being simply open to the environment in a further sense, a sense which concerns the nature of visual experience rather than its epistemic role. Roughly: to have a visual experience with a certain phenomenology or qualitative character is just to experience or represent certain aspects of the visible environment. I'll start by making this more precise.

Broadly following Michael Tye (1995), the proposal is often expressed as:

Intentionalism

> To have a visual experience with a certain qualitative character is just to have an experience with a certain intentional content.

An intentional content is an abstract item to which a subject is related insofar as she represents how things are. If you represent your environment as instantiating *P*, you are *ipso facto* related to an intentional content which attributes *P* to your environment. Intentional contents are often taken to be propositions, though this is not mandatory (Crane 2009). According to *Intentionalism*, to have an experience with a certain qualitative character is just to represent things as being a certain way.[72]

We can express a proposal in the same spirit without assuming that visual experience is a form of representation. For example, John Campbell (2002) understands visual experience as the 'simple relation' of experiencing physical objects and their properties – a relation which cannot be analysed in terms of veridical visual representation. Accordingly, we can specify openness to the visible environment as:

Simple Relationism

> To have a visual experience with a certain qualitative character is just to bear the experiential relation to certain particulars and properties of the visible environment.[73]

To get to grips with these proposals, it will help to consider a structural feature which they share. G.E. Moore claimed that there is a 'common element, … "consciousness", in respect of which all sensations are alike'. He said that sensations or conscious experiences vary only in a second, 'distinct' element: 'the object of consciousness' (1903: 444). Of the two hypotheses above, *Simple Relationism* is the more faithful to the detail of Moore's discussion (Campbell 2009). But we can exploit the abstract claim I quoted to understand what *Intentionalism* and *Simple Relationism*

[72] Tye says that qualitative character is identical with intentional content. As I formulate it, *Intentionalism* does not require this. For example, we might treat as primitive the identification: a subject's *having* an experience with a certain qualitative character is her having an experience with a certain intentional content. I prefer *Intentionalism*, because it doesn't commit us to saying that qualitative characters are abstract items. Similar comments apply, *mutatis mutandis*, to the formulation of *Simple Relationism* below: strictly speaking, *Simple Relationism* does not entail that qualitative characters consist in aspects of the environment.

[73] So formulated, *Simple Relationism* allows that the sheer identity of a seen object, rather than any of its visible properties, might contribute to the qualitative character of visual experience. I don't know whether that possibility is genuine, but nothing here turns on it. I'll be concerned only with the experience of properties.

have in common. If we treat Moore's 'common element' as a two-place relation – call it '*R*' – then we can characterise openness as follows:

Relationism

> To have a visual experience with a certain qualitative character is just to bear *R* to a certain item.

It can be tempting to think of consciousness as a kind of inner glow, a monadic property of the mind or brain. *Relationism* denies this. It denies that having a visual experience with a certain qualitative character is, in part or whole, a matter of the subject's instantiating monadic properties. According to *Intentionalism*, *R* is a relation between a subject and an intentional content. Where intentional contents are propositions, *R* is roughly the relation*: being conscious that.* According to *Simple Relationism*, *R* is a relation between a subject and aspects of her environment – *R* is roughly the relation: *being conscious of.*[74]

By itself, *Relationism* does not entail openness to the visible environment in particular. One traditional version is:

Sense Data

> To have a visual experience with a certain qualitative character is just to bear *R* to certain sense data.

Sense Data characterises visual experience as openness to sense data. *Relationism* characterises openness to the visible environment only once we specify certain *relata* for *R*. *Simple Relationism* achieves this by specifying the particulars and properties of the environment to which we are open. *Intentionalism* achieves it by specifying intentional contents

[74] While Campbell 2002 characterises visual experience as a two-place relation between subject and visible scene, Campbell 2009 characterises it as a three-place relation between subject, standpoint and visible scene. We need not worry about that for present purposes. Since *R* is a two-place relation, *Simple Relationism* must have it that the object of visual experience is a complex item: a particular's instantiating certain properties. That one term of the relation is complex in this way generates some interesting challenges for the theory, discussed by Jacob (ms), but this complexity is consistent with the view that *R* is a two-place relation.

which consist in nothing more than an attribution of certain particulars and properties to the environment.

(In this respect, we can contrast intentional contents not only with sense data, but also with sentences. Like intentional contents, sentences about the environment essentially attribute properties to the environment. But unlike an intentional content, a sentence has further features – its constituent words and syntax – of which you must be aware in order to be aware of what the sentence attributes. For this reason, if you can bear *R* to a sentence about the environment, doing so does not constitute openness to the environment.)

Presumably there are further ways of specifying *Relationism* which also characterise openness to the visible environment. I don't want to prejudge whether any of them might make a difference to the issues discussed here. But for present purposes we can understand openness to the visible environment as follows:

Phenomenal Openness

> Either *Intentionalism* or *Simple Relationism*.

Philosophers sometimes argue that we must accept *Phenomenal Openness* if we're to do justice to the way in which visual experience explains knowledge of the environment (Campbell 2002; Pautz 2010). More often, they appeal to another comment of Moore's. He noted that when you reflect introspectively on visual experience, you become aware only of the 'object of consciousness', rather than the 'common element'; the latter 'seems, if I may use a metaphor, to be transparent'. For example, if you introspect a visual experience of something's blueness, you 'look through it and see nothing but the blue' (1903: 446).

As we saw in Chapter 5, we should be wary of over-generalising from Moore's observation. For example, we found no motivation for:

Transparency

> A subject knows through introspection that her conscious visual representations have a property *P*, only if the objects of those representations seem to her to have *P*.

In particular, we saw that it's compatible with Moore's observation that *differences* among visual experiences should be introspectible, without differences appearing to obtain in the environment. Moore himself qualified his observation, noting that 'awareness', 'the other element' of consciousness, '*can* be distinguished if we look attentively enough, and if we know that there is something to look for' (1903: 450). On the other hand, it's striking that it does seem possible, in principle, exhaustively to describe the introspectible character of your visual experience just by describing the scene you see. One way to accommodate these apparently conflicting points is to take the description to be exhaustive only in the following sense: every introspectible fact about your visual experience is deducible from facts about which scene you experience or represent. For example, if two experiences differ in the properties they represent, this difference is deducible from the first-order facts about what the experiences represent. And the fact that you are conscious – the fact that the 'common element' is in place – is deducible just from the fact that the scene forms an object of introspection.

Some philosophers are sceptical about even this limited commitment to transparency, and I won't pursue it here. For our purposes, it will suffice that many theorists have taken Moore's observation to be obviously correct. Even those who ultimately disagree with Moore must admit that his observation has proved compelling.[75] That itself is an important *explanandum* for theories of visual experience, and *Relationism* is well-placed to explain it. If having a visual experience with a certain qualitative character does not consist in instantiating monadic properties, this explains why no such properties seem to be apparent to introspection (Harman 1990). More positively, if having a visual experience with a certain qualitative character is standing in the experiential relation to certain particulars and properties of the environment, as *Simple Relationism* has it, that explains why introspection seems to reveal these particulars and properties: their presence in the

[75] In arguing for skepticism here, Kind notes the 'widespread agreement that experience is transparent' (2003: 228), where transparency entails that it's impossible to introspect intrinsic qualities of visual experience. Kind proposes an error theory: philosophers have mistaken the fact that it's *difficult* to introspect intrinsic qualities of experience for transparency in this stronger sense. This, she suggests, is truer to Moore's discussion. However, Moore makes clear that the 'common element' does not distinguish one experience from another. So it's hard to see what sort of intrinsic qualities this could include. Moore's 'common element' is more naturally construed as a relation.

environment is constitutive of your having the experience which you introspect (Campbell 2009).

This latter, positive explanation is a little murkier for proponents of *Intentionalism* (and murkier than they often recognise). According to *Intentionalism*, *R* is not a relation to particulars and properties of the environment, but a relation to an intentional content which attributes particulars and properties to the environment. Yet Moore's observation is that introspection reveals properties of the environment, not that introspection reveals intentional contents which merely attribute those properties. So to explain Moore's observation we need some account of why, in taking a relation to a content as the object of introspection, you automatically seem to take the content's *subject-matter* as the object of introspection. To be sure, there is nothing more to the content than an attribution of certain particulars and properties to the environment. But it doesn't obviously follow that what's attributed, rather than the attribution thereof, should be what introspection seems to reveal. However, I won't press this issue here. It does not seem to me implausible that defenders of *Intentionalism* could fill in this missing step.[76]

Relationism is also theoretically attractive because it holds the promise that we might explain the qualitative character of visual experience in terms of information-gathering by the visual system. One sort of explanation here is reductive, and proceeds by appeal to *Intentionalism*. Having a visual experience with a certain qualitative character is identified with having an experience with a certain intentional content. And having an experience with a certain intentional content is identified, in turn, with being in a neural state which bears information in a certain way. For example it's argued that to represent ϕ is be in a neural state which depends asymmetrically on ϕ (Fodor 1990), or to be in a neural state which has the biological function of indicating ϕ (Dretske

[76] Tye says the explanation is 'that visual phenomenal character is representational content of a certain sort, content into which external qualities enter' (2002: 141). I don't see that this helps. Tye glosses 'enter' as follows: 'Qualities entering into the content are qualities the world (or things within the world) seem to the subject of experience to possess' (150). This leaves open the broadly Russellian view that properties are constituents of intentional contents, as well as the broadly Fregean view that intentional contents merely denote properties without having them as constituents. But even on the Russellian reading, it's not obvious why introspection of a relation to a proposition should 'look through' to the proposition's constituent properties in Moore's sense. For example, it would be surprising if you could in this sense look through to uninstantiated properties by reflecting on propositions about them.

1995). So having a visual experience with a certain qualitative character is just being in a neural state which meets certain causal and counterfactual conditions – conditions which may be understood in the terms of the natural sciences.

It's not obvious that *Relationism* allows for a reductive theory of visual experience. Moore claimed that materialism is tempting only because the 'common element' tends to escape our notice when we introspect (1903: 446). Franz Brentano (1874) argued that the intentional relation is irreducible to any of the phenomena studied in the natural sciences. More recently, Adam Pautz (2009) has argued that *Relationism* is inconsistent with any such reductive theory of consciousness. But even if that's right, *Phenomenal Openness* holds the promise of a less ambitious explanation.

Suppose that *R* cannot be understood wholly in other terms: to specify *R* we have to use the notion of conscious experience. Still, we may understand what differentiates one visual experience from another in terms of the different *relata* for *R*. For example, if *Simple Relationism* is true, then we will understand what it is to have a visual experience with one qualitative character rather than another once we understand what it is for the environment to contain certain particulars and properties rather than others. If *Intentionalism* is true, we will understand what it is to have an experience with one qualitative character rather than another once we understand what it is for intentional contents to differ in the relevant ways. (Of course, it may not be straightforward to understand either of these things: a constitutive account of the colours, for example, may be no more straightforward than a constitutive account of intentional contents.)

Phenomenal Openness also provides for satisfying causal explanations of visual experience, in addition to this explanation of what it is to have one visual experience rather than another. Where visual experience is veridical, your bearing *R* to one item rather than another is causally explained in terms of the information which your visual system gathers from your environment – how else could your experience gets things right? Since having an experience with one qualitative character rather than another is just bearing *R* to one item rather than another, we can explain why your experience has one qualitative character rather than another by reference to the information-gathering carried out by your visual system. If *Phenomenal Openness* is true, this causal explanation is satisfying in that there are no further, monadic properties of visual experience which it leaves explained. In principle at least, this gives us

some traction in the experimental study of visual consciousness. The variable to manipulate, if you want to manipulate visual consciousness, is just the information-gathering carried out by the visual system.

The weaker, non-reductive explanations just described identify having a visual experience with a certain qualitative character with bearing *R* to a certain item. The weakening lies in their refusal to identify bearing *R* to a certain item, in turn, with meeting a causal or counterfactual condition. The explanation is sometimes weakened at the first stage instead. *Intentionalism* is weakened to a supervenience thesis, rather than an identity thesis (e.g. Tye 2002; Speaks 2010). For the more abstract *Relationism*, this gives us:

Supervenience Relationism

> Necessarily, visual experiences *E* and *F* differ in qualitative character only if: the item to which a subject bears *R* in having *E* differs the item to which a subject bears *R* in having *F*.

Supervenience Relationism says that the qualitative character of a visual experience supervenes on the items to which you're related, in having the experience. This is entailed by, but does not entail, the identity hypothesis *Relationism*. *Supervenience Relationism* does not tell us what it is to have a visual experience with a certain qualitative character. Rather, it tells us what it takes for experiences to differ in their qualitative character, or what settles the facts about how visual experiences differ. Again, this provides for an explanation of the qualitative character of visual experience, though not a reductive explanation. As before, we could further specify *Supervenience Relationism* to characterise openness to the environment, and we could specify this still further in terms of either a simple experiential relation or an intentional relation.

I'm going to focus on a challenge facing the defender of *Phenomenal Openness*. But it should be obvious that this is also a challenge for the defender of a theory which weakens *Phenomenal Openness* to a supervenience thesis.[77]

---

[77] When *Intentionalism* is weakened to a supervenience thesis, it's even trickier to explain Moore's observation about transparency. Why should *A*'s supervening on *B* ensure that awareness of *A* is awareness of *B*? Instead of supervenience, we might here exploit the notion of a constitutive (minimal sufficient) condition explained in §3.1. The two notions are not equivalent, and minimal sufficient conditions have some advantages if we want

6.2 The Challenge from Attention

Philosophers have raised various counterexamples to *Phenomenal Openness*, and its defenders have offered various responses. The counterexamples involve a difference between the qualitative character of two experiences, which – it's urged – is not accompanied by a difference in the experiences' intentional content, or by a difference in the particulars and properties experienced. The responses attempt to show that there is, in fact, such an accompanying difference. I won't review them here.[78] I'll be concerned specifically with the challenge generated by conscious attention, which I think runs especially deep, as discussed in Chapter 1.

Figure 22

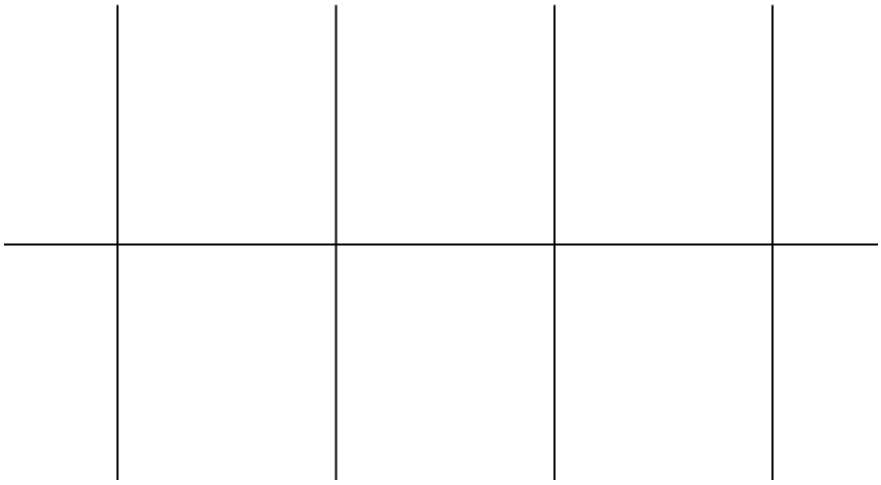

to capture the idea of that which *realises* qualitative character. *A* may supervene on *B* without *B*'s being any part of a minimal sufficient condition on *A*. For example, your visual state supervenes on the entire state of the universe. Intuitively, the entire state of the universe does not realise your current visual state; much of the universe is idle in that respect. Equally, *B* may be a minimal sufficient condition on *A* without *A*'s supervening on *B.* Suppose that each of *B* and *C*, severally, forms a minimal sufficient condition on *A.* Then a difference in *A* entails neither a difference in *B* nor a difference in *C.* Yet intuitively, I think, each of *B* and *C* realises *A.* I won't press this issue here. The problems I'll discuss for *Phenomenal Openness* are equally problems for theses which weaken *Phenomenal Openness* in terms of minimal sufficient conditions on having a visual experience with a certain qualitative character.

[78] For some counterexamples, see e.g. Peacocke 1990 and Nickel 2007. For some responses, Tye 2002 and Nanay 2010.

Jeff Speaks sets out the challenge by reference to Figure 22. Hold the figure so that it occupies more or less the whole of your field of view. If you attend to the intersection of the horizontal line and the second vertical line from the left, your visual experience seems to differ in phenomenal character from the experience you enjoy if you attend to the intersection of the horizontal line and the second vertical line from the right. Speaks claims that the experiences so differ, but are identical in intentional content; so *Intentionalism* is false. Similarly, if the experiences so differ, but are identical in the particulars and properties which you experience in having them, then *Simple Relationism* is false.

Why should we accept that the two experiences do not differ in their intentional contents, or in the particulars and properties experienced? Nothing in the scene seems to differ when you shift attention as directed. But we've seen that this is compatible with both dilation and translation in the determinable properties which visual experience represents.[79]

Bence Nanay (2010) defends *Intentionalism* against Speaks' proposed counterexample by claiming, in effect, that there is dilation in the properties experienced: he says that each experience represents more determinate properties than the other at the part of the scene that's attended. Against this line of response, Speaks argues that 'given the simplicity of the figure, it does not seem plausible to claim that one experience represents a given portion of the lines with more detail or determinacy' (2010: 329). But we've seen that *every* visual experience represents determinable spatial properties, and so could, in principle, come to represent more determinate spatial properties.

However, this is not an issue to be settled completely *a priori*. We can employ Yeshurun and Carrasco's empirical test (§4.3) to assess whether it's plausible that the phenomenal change in Speaks' case consists in a dilation in the properties experienced. Yeshurun and Carrasco (1998) found that involuntary covert attention increases visual spatial resolution. It follows, I have argued, that this form of attention involves dilation in the properties which visual experience represents. Now in Speaks' case, attention is covert: the figure occupies more or less the whole of your field

[79] As in previous chapters, I'll use the expression 'represents *P*' loosely, to cover visual experience in which a property figures essentially whether or not this is construed, strictly speaking, in terms of visual representation. For instance, I do not here assume that *Intentionalism* rather than *Simple Relationism* is the view to be defended.

of view, so you cannot shift fixation overtly from one intersection to the other. But attention here is voluntary not involuntary: you shift your attention on purpose, deliberately following the instructions. And there is a growing body of evidence that voluntary covert attention has systematically different effects from involuntary covert attention (Prinzmetal and Landau 2008).

Yeshurun et al. (2008) tested the hypothesis that voluntary covert attention has the same effect as involuntary covert attention, similarly increasing the visual spatial resolution. They used a similar experimental procedure, except that instead of cuing attention by flashing a small bar, they used a semantic cue. Before presenting a texture-segregation stimulus (Figure 23), they showed subjects numbers that indicated which part of the scene they should attend to, and allowed them enough time to attend voluntarily. They then repeated the procedure, before asking subjects to report which of the two stimuli contained the unique texture.

Figure 23

A texture-segregation stimulus

Recall that, for some stimuli, texture-segregation is harder where spatial resolution is too high. As a result, involuntary covert attention either impairs or improves performance at the task, depending on the location of the stimulus relative to the fovea. Where this location is such that spatial resolution is at the right level without attention being cued, cuing involuntary covert attention makes subjects worse at the task (Yeshurun and Carrasco 1998). By contrast, Yeshurun et al. (2008) found that voluntary covert attention only ever improved performance. Since representing more determinate spatial properties ought to inhibit performance at some locations, we can't explain the improved performance just in terms of the idea that voluntary covert attention makes visual experience represent more determinate spatial properties.

The test does not confirm Nanay's claim that, in a case like Speaks', attention makes visual experience represent more determinate properties.

Yeshurun et al. suggest two possible interpretations, to which I'll add a third:

(i) Voluntary attention reduces the noise in early visual processes, without altering visual spatial resolution.

As modelled in signal detection theory, a reduction in noise may make a system more reliable in selecting among possible representations, without altering the semantic value of the representations generated (MacMillan and Creelman 2005). If (i) is correct, the phenomenal change brought about by voluntary covert attention does not entail visual experience's representing more determinate spatial properties.

(ii) Voluntary attention is more 'flexible' than involuntary attention, in that voluntary attention either increases or decreases visual spatial resolution to suit the task at hand.

This looks more promising for the defender of *Phenomenal Openness*, but in fact it should just prompt a refinement of the putative counterexample. Take a case in which, without focussing attention on it, you're already seeing a stimulus in just the spatial resolution required for your task. Now focus attention on the stimulus. The determinacy of the spatial properties represented ought to remain fixed. Still, it's hard to believe that under these circumstances attending will bring about no phenomenal change – that conscious attention is simply impossible. So here we should expect a phenomenal change which does not entail dilation in the spatial properties experienced. Against this, you might insist that such cases involve no phenomenal change. But that would be a surprising empirical result, so this strategy inherits a substantial burden of empirical proof.

(iii) Voluntary attention increases spatial resolution just as involuntary attention does, but voluntary attention in some other way ensures that subjects are better at detecting the unique texture, mitigating the detrimental effects of resolution that's too high.

I'll assume that, if attention in some other way ensures that subjects are better at detecting the unique texture, this positive effect of voluntary attention has some phenomenology or phenomenal correlate. One might in principle deny this. But again this would be a surprising result. Covert voluntary attention seems, on the face of it, to be a paradigm case of improved performance *through* heightened consciousness. So we should expect there to be a phenomenal difference in virtue of which voluntary attention has its positive effect on performance at the task. According to (iii), this difference does not involve dilation in the spatial properties which visual experience represents.

In principle again, the phenomenal differences brought about by attention might entail differences in what visual experience represents other than dilations in the spatial properties experienced. For instance voluntary attention might lead to a translation, rather than a dilation, in the spatial properties experienced: by making the unique texture appear larger in the way explored in Chapter 5, attention might make the unique texture more visible. Or attention might lead to dilation with respect to non-spatial properties of the unique texture, such as its colour. However, there is no such effect which is both empirically plausible and apt to explain subjects' improved performance at the task. For example, the evidence suggests that voluntary covert attention does *not* lead to translation in the spatial properties which visual experience represents: voluntary attention does not lead subjects to overestimate spatial properties the way in that involuntary attention does (Prinzmetal et al. 1998; Gobell and Carrasco 2005).[80] And the claim that voluntary attention makes visual experience represent more determinate colours does nothing to explain subjects' improved performance. So there is no clear candidate for a difference in what visual experience represents, which could constitute the phenomenal effect of attention.

Interpretations (ii) and (iii) draw on the evidence about involuntary attention, claiming that voluntary attention has somewhat similar effects.

[80] Liu et al. (2009) found that voluntary attention has an effect similar to that of involuntary attention on the visual experience of brightness contrast: if we interpret their results in the manner of Chapter 5, involuntary attention leads to a translation in the contrast experienced. It does seem plausible that increased apparent contrast should make a unique texture easier to detect accurately. However, to that extent increased apparent contrast should lead to an increase in the resolution of vision. And increased resolution should *impair* performance at some tasks. By contrast, the effect we're looking for only ever improves performance.

Interpretation (i) denies this. But whichever interpretation we prefer, the evidence favours the view that, although attention sometimes makes visual experience represent more determinate properties, covert voluntary attention to what you see also brings about phenomenal differences which entail no difference in what visual experience represents.

6.3 Attention and Modes of Consciousness

In Chapter 4 I distinguished between *content focalization* and *mode focalization* of consciousness. Content focalization is a change or variation in what experience represents. Mode focalization is a change or variation in the mode of consciousness by which what's represented is represented.

In §6.2 we identified phenomenal differences between experiences of attentive vision which – the evidence suggests – entail no difference in what visual experience represents. In this section, I'll look at some ways of understanding these differences in terms of the mode focalization of consciousness, and inconsistently with *Phenomenal Openness.* Then in the remainder of the chapter I'll argue for an alternative approach, and assess how this alternative sits with respect to *Phenomenal Openness*.

We've already seen some ideas about how attention might involve differences in the mode of consciousness by which things are represented. Ned Block (2010; §5.3 above) claims that visual experience consists of a representational content together with a mode of presentation of that content, such that the mode of presentation of an experience fixes its phenomenology, making the object of experience appear a certain way. For example, the mode of presentation of a subject's experience fixes a certain 'perceived gap size', 'perceived contrast' and so on. Yet the mode of presentation of an experience may vary independently of its representational content. Where the mode of presentation fixes a perceived gap size of 0.23°, visual experience need not represent the gap as 0.23° wide. In that sense, visual experience is characterised by 'mental paint', a way things appear to a subject which is not a matter of how she represents them to be. Block claims that some effects of attention consist in changes in the mode of presentation of an experience which entail no change in its representational content.

Block's proposal is inconsistent with *Phenomenal Openness*: to have a visual experience with a certain qualitative character is not just to be related to certain aspects of the environment, or to a certain intentional

content; it's in addition to enjoy a certain mode of presentation of that content. However, the proposal does not help to characterise the phenomenal differences identified in §6.2. These differences are effected by *voluntary* attention. Block introduced the proposal in order to explain Carrasco's data about mistaken visual judgement, and those data concern involuntary attention.[81] The evidence suggests that voluntary attention works quite differently: it does not lead to parallel mistaken judgements about colours and shapes (Blaser et al. 1999; Prinzmetal et al. 1998), and it improves performance at Yeshurun et al.'s (2008) task.

Block's account of modes of presentation is unusual, in that he understands them by reference to 'perceived' properties which the objects of experience might (in principle) have – for example a perceived gap size of 0.23°. More typically, modes of presentation are understood as ways in which a content may be presented which cannot be indexed to possible properties of the environment; possible properties of the environment figure in the content of experience, not in the mode of presentation of that content.

Sebastian Watzl (2011) claims that all conscious experiences admit of a systematic ordering in terms of how focal or peripheral they are within the stream of consciousness. That is, experiences can be ordered according to the extent to which their subject-matter is a focus of attention: necessarily, wherever a subject is conscious of ϕ, ϕ is to some degree focal or peripheral in her stream of consciousness; if she is simultaneously conscious of something else, χ, then ϕ and χ admit of a principled ordering in this respect. Watzl argues that the ordering is primitive: the facts about how focal an experience is cannot be reduced to facts about the determinacy of experience, say.[82] Unlike Block's 'mental paint', these differences in mode of presentation are apt to capture the idea that you may be *more conscious* of something, in a way which entails no difference in what your experience represents.

Watzl's proposal is straightforwardly inconsistent with *Phenomenal Openness*. According to Watzl, bearing the generic conscious relation *R* to an intentional content, or to certain particulars and properties of the environment, is never sufficient for having visual experience with a certain qualitative character. In addition to the facts about these *relata* for

[81] The exception is Liu et al. 2009. See footnote 80.

[82] Watzl's argument here appeals to *Reproduction*, as criticized in §4.3.

R, there are the facts about how focal or peripheral the objects of experience are, within the stream of consciousness.

As a result, even if bearing *R* to a certain item is reducible to a causal or counterfactual condition, having a visual experience with a certain qualitative character is not reducible to that condition. And the less ambitious, non-reductive explanations of the qualitative character of visual experience are similarly compromised. We cannot give a constitutive account of the differences between visual experiences just in terms of the different intentional contents, or the different aspects of the environment, to which subjects are related in having those experiences. Nor can we give a satisfying causal explanation of why someone enjoys the visual experience she does, just in terms of the information which her visual system gathers from the environment.

However, the data don't require that we accept a position as strong as Watzl's. He assumes that ordering in the relevant respect is essential to every episode of visual experience. An argument for this claim would be beset by problems of the kind explained in Chapter 3. Suppose, for the sake of argument, that every instance of visual experience for which we have first-personal or third-personal evidence admits of ordering in the way explained: it is relatively focal or peripheral within the stream of consciousness. This might reflect a condition on the evidence being in place, rather than a condition on visual experience itself. The explanation *might* be that every instance of visual experience is essentially ordered in the relevant way. But equally, what imposes the ordering on visual experience might be attending in the way required for there to be evidence of visual experience. That is, it might be that you can introspect or report visual experience only by attending to what you see in a way which makes it relatively focal or peripheral within your overall experience.

Speaks' own response to his counterexample is more modest, in that it requires only a contingent, computational connection between visual experience and attention. Speaks suggests that we treat attention as a *sui generis* modality of conscious experience, analogous to but constitutively independent of visual experience. He proposes that attention is an additional, distinctive way of being related to the same items to which one is related in having visual experience (2010; §4.2 above). We can understand the mode focalization of consciousness in these terms: in addition to enjoying visual experience of a scene, you are conscious of a scene in a further way; you experience the scene in more

than one mode of consciousness. Although Speaks does not emphasise this, his proposal has the right form to explain why we might be better at identifying things when we attend to them – for instance, why Yeshurun et al.'s subjects were better at identifying the stimulus with a unique texture in their texture-segregation task. For example, Declan Smithies agrees that attention is a 'distinctive mode of consciousness', and argues that to attend to ϕ is to select ϕ in consciousness, in a way which makes information about ϕ available 'for use in the rational control of action, reasoning and verbal report' (2011: 257; §4.2).

In Chapter 4 I rejected these theories of attention as an essentially non-visual, *sui generis* mode of consciousness, on the grounds that attention sometimes takes the form of content focalization within visual experience. However, that's consistent with the view that attention *sometimes* takes the form of a *sui generis* mode of consciousness. Nothing I've said rules out the possibility that such a mode of consciousness explains subjects' performance in Yeshurun et al.'s task, or even the possibility that such a mode of consciousness is typically one component in the phenomenology of voluntary attention.

One qualm we might have about Speaks' proposal is as follows. Because he sets out the challenge to *Intentionalism* by reference to a pair of cases, Speaks casts attention as a binary phenomenon: either you are conscious of ϕ in the modality of attention, or you are not so conscious of ϕ. But on the face of it, attention is a matter of degree: you can attend to something more or less. This is a standard assumption in the science of attention. For example attention is treated as spreading, to a limited degree, to objects near a focus of attention (Mack and Rock 1998; §3.2 above). It's also a natural assumption from the point of view of the phenomenology of attention. For example, suppose that while you attend to one of the intersections in Speaks' grid, you're distracted by something moving at the edge of your field of view. On the face of it you might continue attending to the intersection, but do so less fully than when not distracted.

However, we can understand Speaks' proposal consistently with this feature of attention. We can understand attention as admitting of degree in Watzl's sense, except that here the ordering is limited to consciousness in the modality of attention: necessarily, wherever a subject is conscious of ϕ in the modality of attention, ϕ is to some degree focal or peripheral in her stream of consciousness; if she is simultaneously conscious of something else, χ, in the modality of attention, then ϕ and χ

are ordered in this respect. This leaves it open whether or not *all* experiences admit of ordering in the way explained.

With or without this addendum, Speaks' proposal is inconsistent with *Phenomenal Openness*, since as we saw *Phenomenal Openness* entails:

Relationism

> To have a visual experience with a certain qualitative character is just to bear *R* to a certain item.

Speaks' proposal entails that bearing the generic conscious relation *R* to a certain item is insufficient for having a visual experience with a certain qualitative character. In bearing *R* to an item, you might enjoy *sui generis* attentive experience rather than visual experience. So bearing *R* to a certain item is never specific enough to form a sufficient condition on visual experience. A sufficient condition requires a more specific conscious relation, as well as more specific *relata*.

We can compare this objection to *Phenomenal Openness* with an objection from non-visual perceptual experience. Suppose you think that two experiences of a certain shape may differ, just insofar as one is a visual experience and the other is an experience of touch. This is inconsistent with *Phenomenal Openness*: if experiences may differ in this way, then in bearing the generic conscious relation *R* to the shape (or to an intentional content about the shape) you might enjoy an experience of touch, rather than a visual experience. *Phenomenal Openness* requires that the modality of perception does not make a phenomenal difference in this way; the differences between visual and haptic experience must lie in what they represent – for example in the fact that only visual experience represents colours. Similarly, *Phenomenal Openness* requires that the differences between attentive and inattentive experience lie in what they represent, not in the mode of consciousness by which it is represented.

Speaks argues that *Intentionalism* should be amended to accommodate his *sui generis* modality of attention. For the more general *Relationism*, we can do this in terms of conscious relations more specific than the generic conscious relation *R*. Call '*P*' the relation you bear to an item in having inattentive perceptual experience. We can then propose:

Intramodal Relationism

> To have a visual experience with a certain qualitative character is just to bear *P* to a certain item.[83]

If we're interested in explaining the qualitative character of visual experience, *Intramodal Relationism* is a significant concession. No reductive explanation can proceed along the lines sketched at the beginning of this chapter. Standing in the intentional relation to a content is not sufficient for having visual experience, rather than experience in a different mode or modality. So even if the intentional relation can be reduced to a causal or counterfactual relation, the more specific, distinctive qualitative character of visual experience will remain unexplained.

Intramodal Relationism is also a significant concession from the point of view of less ambitious, non-reductive explanations. A causal explanation of why your experience has one qualitative character rather than another cannot be given just in terms of the information-gathering carried out by your visual system: even once we know what information has been gathered, what makes an experience distinctively visual will remain explained. Similarly, a constitutive account of what it is to have an experience with one qualitative character rather than another cannot be given just in terms of the different particulars and properties which an environment may contain, or just in terms differences among intentional contents.

However, we need not accept any of the proposals canvassed in this section, or the amendments to *Phenomenal Openness* which they force. Instead, we can understand the phenomenal differences identified in §6.2 as differences in conscious cognition. This approach has the advantage of parsimony. In particular, it makes postulating *sui generis* phenomenal facts about attention otiose. Furthermore, it entails that the phenomenology of attention poses no novel or distinctive challenge to *Phenomenal Openness.* Any limitations we impose on *Phenomenal Openness* are in any case

[83] An alternative would be to replace *R* with a three-place relation between a subject, a mode of consciousness, and a further item (e.g. an intentional content or an aspect of the environment). For our purposes the choice won't matter. Speaks himself proposes a limited supervenience thesis: 'Necessarily, if two perceptual experiences of the same sense modality differ in phenomenology, then they differ in content' (2010: 325). The limitation to differences within a modality achieves an effect parallel to that of *Intramodal Relationism*.

required in an account of the conscious mind which accommodates both conscious perception and conscious cognition.

6.4 Attention and Cognition

In Speaks' case, as in most others discussed in the recent philosophical literature, there is an intuitive sense in which we're dealing with *perceptual* attention. At a minimum, these are cases in which perceptual experience and attention share an object. In looking at ϕ, for example, you both see ϕ and attend to ϕ.

It bears emphasis that, as the term 'attention' is used in English, attention need not meet even this minimal condition. You might say, for example, that a certain abstract problem received a lot of attention in the Eighteenth Century, or that the fiscal crisis occupies the President's attention. Consider MGF Martin's comment:

> What are the most obvious generalizations about attention and thought that form part of the manifest image of these aspects of mind? When I think about the level of subsidy for arable land in the Common Agricultural Policy, I thereby attend to European farming policy. In general, whatever we are prepared to call an object of thought – be it the things thought about, what one thinks about them, or the proposition one thinks in thinking these things – we can also take to be an object of attention. Conscious, active thought is simply a mode of attending to the subject matter of such thoughts.
>
> Martin 1997: 77

According to one traditional proposal, attention is *always* a cognitive phenomenon in this way, even in cases of perceptual attention. Peter Geach claimed that, on the use of 'attend' relevant to epistemology and the theory of meaning,

> "attended to" is a mere word for the … relation of judgement to sense-perception … – "James judges, attending to such-and-such sense-perceptions", meaning nothing more than that his judgement *refers to* what he thus perceives.
>
> Geach 1957: 64

According to Geach, to be attended in this sense is just to be the perceptible subject-matter of a judgement; even in cases of perceptual attention, to attend to ϕ is simply to take ϕ as the subject-matter of a thought or judgement.

Geach was not explicitly concerned with the phenomenology of attention, but his comment is not neutral with respect to it. When we say that *S* attends to ϕ, where ϕ is an object of perception for *S*, we do ascribe a certain conscious experience. According to Geach, what we ascribe here is a thought or judgment which takes ϕ as its subject-matter. We can exploit this idea in an analysis of the phenomenal differences identified in §6.2.

Figure 24

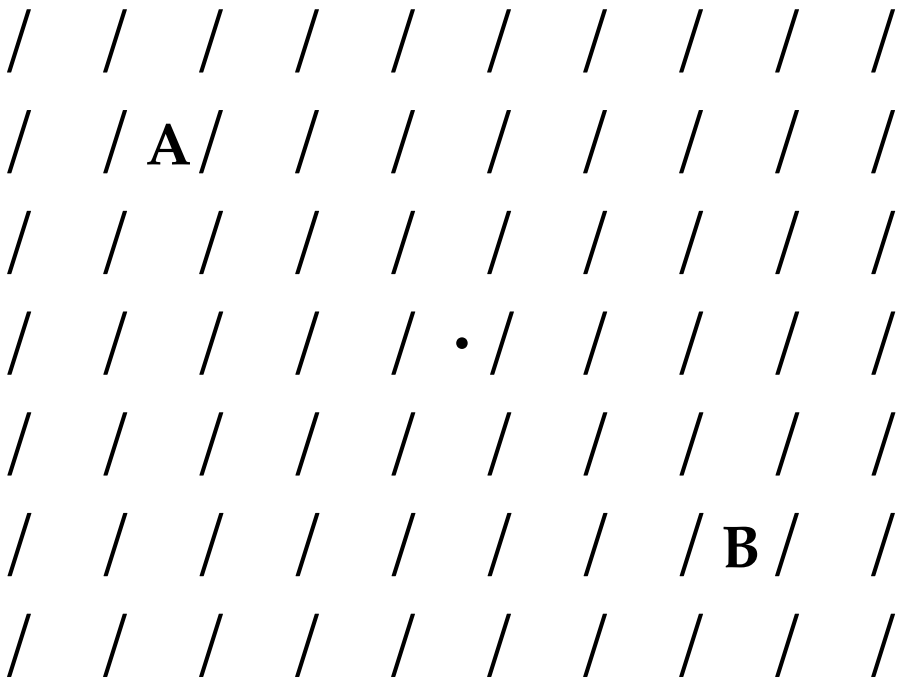

Consider Figure 24. Hold it so that it occupies more or less the whole of your field of view, and attend first to the 'A' then to the 'B'. Call your experience when attending to the 'A' the 'A-Experience'. Call your experience when attending to the 'B' the 'B-Experience'. As I have argued, the evidence suggests that there is a phenomenal or qualitative difference

between the A-Experience and the B-Experience which is not a difference in what visual experience represents. According to the analysis I have in mind, this difference consists in the following: each experience involves a conscious thought which takes a different part of the figure as its subject-matter.

More generally, where attending voluntarily to what you see makes a phenomenal difference over and above any differences in what visual experience presents or represents, this is a difference in conscious cognition. This is just the broadly Jamesian analysis which I introduced in Chapter 4: an episode of conscious attention consists in a 'focalization of consciousness', a change or variation which may occur in visual experience, in conscious cognition, or in both. The Jamesian analysis is a competitor to the accounts of conscious attention canvassed in §6.3: on the Jamesian analysis, we need not postulate a *sui generis* attentive modality of consciousness, or irreducible facts about an attentive mode of presentation of experience. First I'll respond to two objections to the Jamesian analysis. Then in §6.5 I'll explain some of its advantages.

One objection focuses on the logical form of thought and attention. Thoughts, it may be argued, are necessarily propositional. You are not thinking anything, or thinking about anything, unless you're thinking that something is the case. As Martin's comment makes clear, there are episodes of attention which likewise take propositions as their subject-matter. But is this true of every episode of conscious attention which the my analysis purports to capture in terms of conscious cognition? For example, does the difference between the A-Experience and the B-Experience really entail your thinking that something is the case? Do you really entertain some proposition just in attending to the 'A'?

Clearly this raises deep questions about the role of propositions in thought and experience, questions which I won't be able to answer in depth here. But I think we can see fairly easily how to respond to the objection in outline. First, it is a difficult question how much propositional structure is required for a thought. For example, if there can be thoughts with the merely feature-placing content '*A*!' (Strawson 1959), then that might be a plausible candidate for what you think in having the A-Experience.

Second, suppose we accept that thoughts necessarily concern propositions which are richer in logical form. It's not hard to find propositional contents that are reasonable candidates for being the contents of attention, even in the simple case of the A-Experience.

According to many philosophers, perceptual experience is itself a relation to complex propositions. For example, insofar as the A-Experience involves perceptual experience, it is experience of a complex proposition which ascribes certain shapes to the various parts of Figure 24, certain orientations to some lines, and so on. To the extent to which this is plausible, presumably it's equally plausible that attention is a relation to some part of this complex proposition – the part which concerns the attended part of the figure.

Moreover, even if perceptual experience is not itself propositional, perceptual experience is at any rate experience of complex states of affairs. For example you see not just the line, but also the line's having a certain orientation. Why should propositions about these states of affairs not form the propositional contents of attention? For example, suppose that conscious attention necessarily predicates properties of particulars. Even if visual experience does not involve this predicational structure, visual experience involves a structure of parallel complexity. That much is implied just in a line's visibly having a certain orientation.[84]

But what, specifically, is the logical form of the supposed propositional contents of conscious attention? Geach's own proposal is objectionable, from the point of view of the phenomenology of attention. He took the sort of judgement involved in perceptual attention to be the same in form as judgements made in non-perceptual contexts:

> [W]hat is the difference between "there are white cats", "some cats are white", on the one hand, and on the other hand "these cats are white"? How do the judgements they express differ? What constitutes the reference to a particular set of cats? In all such cases, I should maintain, there is no difference to be found on the side of the judgement itself. What we may call the intelligible content of the judgement is the same …but an act of judgement performed in a particular sensory context may thereby be referred to particular sensible things.
>
> Geach 1957: 63-4

[84] I rely here on the premise that properties such as orientations are present in visual experience in the sense explained in §2.2. On a standard semantics, '*S* sees ϕ', where ϕ is a complex state of affairs, need not ascribe to *S* an experience with complex structure. See Barwise 1981.

There are familiar reasons for thinking, *pace* Geach, that the proposition expressed by 'these cats are white' differs in form from the proposition expressed by 'some cats are white'. But without exploring the broader issues about demonstrative expressions, in the present context we can focus on a phenomenological concern.

On the face of it, the experience of thinking 'there are some cats', or even the indexical 'there are some cats here', is quite different from the experience of attending to the some cats. The experience of attending to the cats seems to trade on or exploit visual experience of the cats, so that – unlike the experience of thinking 'there are some cats here' – the experience of attending is an experience which you could not have in the absence of the visual experience. Presumably this sort of phenomenological evidence is defeasible, especially when it comes to modal claims about the conditions under which an experience is possible. But if we're hoping to explain the phenomenology of perceptual attention in terms of propositional thought, we have a *prima facie* reason to assume that this thought will reflect, in its form, the apparent dependence of perceptual attention on perceptual experience.

Wayne Wu's (2011a) cognitivist account of conscious attention is designed to accommodate this assumption. In Wu's terminology, the distinctive kind of consciousness which constitutes conscious attention to an object of visual experience is 'phenomenal salience': where one seen object rather than others is an object of conscious attention, this object is 'phenomenally salient'. Wu argues that phenomenal salience consists in visual-demonstrative cognition – cognition which we might express by talking about 'these cats', which differs in form from existentially quantified thought, and which depends essentially on visual experience of its objects. Where an object is phenomenal salient, this consists in the fact that 'the subject is demonstratively cognitively representing that object and not others concurrently perceived' (2011a: 111). So for example the A-Experience differs qualitatively from the B-Experience in that the 'A' rather and not the 'B' is an object of demonstrative cognition.

There are different ways of understanding the distinctive form of demonstrative cognition. For example, we might propose that 'these cats are white', uttered by someone currently seeing some cats, expresses a proposition of which the cats themselves are constituents (Kaplan 1989). Alternatively, we might propose that the utterance expresses a Fregean proposition of which no concrete objects form a part, but which can be grasped only by someone currently seeing the cats (McDowell 1998). For

present purposes we need not choose between these alternatives. What's important is just that someone who thinks a proposition of the relevant form is, necessarily, enjoying visual experience of the objects referred to demonstratively.[85]

In this way, Wu's account captures the apparent dependence of attention on visual experience. To this extent his account captures the phenomenology of attention, and I agree that where attention takes the form of conscious cognition, cognition is likely to be demonstrative in form. However, the phenomenology of attention also has a further feature which is inconsistent with Wu's account: conscious attention sometimes includes a genuinely visual change, a change in visual experience itself, rather than a change in conscious cognition.

According to Wu, episodes of conscious attention and episodes of visual experience are always distinct. The connection is, at most, that episodes of attention require episodes of visual experience:

> Perceptual consciousness [is] essential to phenomenal salience, for the demonstratives at issue are demonstratives that are grounded in ongoing perception, demonstratives that are available precisely because what is demonstrated is consciously perceived. … Demonstrative representations make a difference to what it is like for the subject in part due to this dependence on conscious perception. The phenomenology of demonstrative thought will partly reflect the phenomenology of perception on which it is based.
>
> Wu 2011a: 115

There is therefore a compelling objection to Wu's account. On the face of it, conscious attention sometimes has a closer connection than this with perceptual experience. The phenomenology of attention sometimes seems to include a perceptual change or variation – not merely to depend upon the contents of perceptual experience. And as I argued in Chapter 4, there is decisive empirical evidence that this is in fact the case. Attention sometimes takes the form of dilation in the properties which visual experience represents. Conscious attention is not, then, just a matter of conscious cognition.

---

[85] The questions here are somewhat different from those about demonstrative thought which I discussed in §2.1. Here we may limit our concern to thoughts which are demonstrative with respect to objects. There the focus was on thoughts which are demonstrative with respect to properties.

Unlike Wu's analysis, the analysis I propose is consistent with this finding. On my broadly Jamesian analysis, attention is a focalization of consciousness which may occur either in visual experience or in conscious cognition; in many cases, it occurs in both. On the one hand, where the focalization of consciousness takes a visual form, visual experience represents more determinate properties. On the other hand, in addition to enjoying visual experience of an object, you may also take it as the subject-matter of conscious cognition. This also constitutes a focalization of consciousness, a way of becoming more conscious of what you see.

A different objection to my analysis focuses on the explanatory role of conscious attention. In the passage I quoted, Geach is objecting to Russell's claim that attention to ϕ provides a causal and rational explanation of the capacity to form thoughts about ϕ. According to Geach, attention to ϕ cannot provide this, since to attend to ϕ in the relevant sense is already to form a thought about ϕ. To reverse this, one might object to my analysis of conscious attention on the grounds that conscious attention is, in fact, explanatorily prior to conscious thought. For example we've seen that according to some philosophers, conscious attention to the objects of perception causes and justifies thoughts about those objects; conscious attention is the pre-condition which makes these thoughts possible (Campbell 2002; Smithies 2011).

Note, however, that this objection requires a very specific premise. It requires that the phenomenal differences identified in §6.2 are explanatorily prior to conscious judgement about what you see. For example, it requires that the phenomenal difference between the A-Experience and the B-Experience (over and above any differences in what visual experience represents) is explanatorily prior to conscious judgements about the 'A' and the 'B'. We can deny this while accepting much of what's at stake in saying that attention and visual experience cause and justify judgements about the objects of visual experience. We can accept that *acts* of voluntary attention cause and justify these judgements, while denying that they do so in virtue of their distinctive phenomenology. We can also accept that the phenomenology of conscious *visual* experience causes and justifies judgements about its objects, while denying that the phenomenology distinctive of voluntary attention does the same. I'll take each of these two points in turn.

Consider John Campbell's (2002) view that conscious attention causes and justifies demonstrative judgements about the objects of visual experience. Campbell appeals to evidence about attention and 'binding

parameters' in the visual system. Binding parameters are the principles by which the visual system conjoins information about different features of an object such as its colour and shape, enabling you to experience them as features of one unitary object. Features in different categories like these are processed separately. Under demanding conditions, subjects are sometimes able to report accurately what features are present in a scene, yet unable to report accurately which of them are conjoined in particular objects. For example, if you're presented with a large number of objects under time pressure, you may be able to report accurately that there's something red and that there's something square, yet unable to report accurately whether there's something which is both red and square. Roughly, the visual system conjoins different features where they share a location: the capacity to report conjunctions of features accurately is reliably correlated with the capacity to report accurately where features occur in a scene.[86] The mechanism by which this is achieved is a mechanism of selective attention: unitary objects are encoded one after the other, in a process of serial search that's characteristic of selective attention. This shows up in the fact that increasing the number of objects in a scene impairs reports of feature-conjunctions disproportionately, as compared with reports of single features (Treisman 1988).

Campbell argues that attention also exploits visual binding parameters to play a further role: attention causes and justifies visual-demonstrative thoughts – thoughts which we can express by talking about '*that* object', and which we understand by attending to the object as it figures in visual experience. For example, if I point to the 'A' in Figure 24 and say 'That is black and contains a triangle', you can verify what I've said by attending to the 'A'. In doing so, you exploit the principles by which your visual system treats the colour and shape of the thing as features of one unitary object.

Now none of this requires that the distinctive phenomenology of attending voluntarily to what you see causes or justifies demonstrative thoughts. For example, it does not require that the qualitative difference between the A-Experience and the B-Experience is in this way explanatorily prior to conscious judgement about the 'A' or the 'B'. Rather, this account of demonstrative reference is consistent with the claim that voluntary attention has its distinctive phenomenology only to the extent

[86] Coherent movement is also a factor, along with other Gestalt principles (Driver and Baylis 1989).

to which thought about the objects of visual experience is actually achieved. We can compare voluntary attention here with bodily action. In performing a bodily action, you set in train muscular mechanisms which make it the case that you move your body in a certain way. You act only to the extent to which you actually move, but it's nonetheless true both that you exploit the muscular mechanisms, and that those mechanisms cause you to move. In attending voluntarily, in order to understand a demonstrative thought about the 'A', you set in train the mechanism of selective attention which I described above. I'm proposing that you enjoy the A-Experience only to the extent to which you actually think demonstratively of the 'A'; nonetheless, it may be true both that you exploit the mechanism of selective attention in order to have this thought, and that the mechanism causes and justifies the thought.

Campbell's account does require that demonstrative thought is caused and justified by conscious awareness of objects and their locations, rather than by merely subpersonal mechanisms of spatial attention (2002: 34). Like Russell, he takes conscious awareness of an object to be a pre-condition which makes thought about the object possible. There is also good empirical evidence that conscious awareness of an object's location, rather than merely implicit information-processing about its location, is correlated with the capacity accurately to report conjunctions of visible features. Subjects who are unable to report feature-conjunctions accurately are also unable to report objects' locations accurately, but they do show priming effects for information about the objects' locations (Nissen 1985; Robertson 2004). These subjects process information about an object's location, but that doesn't enable them to verify demonstrative thoughts about the object – for that, conscious experience of the location is required.

However, this does not entail that the phenomenology distinctive of *attention* causes and justifies demonstrative thoughts, or is explanatorily prior to such thoughts. For example, it does not require that the qualitative difference between the A-Experience and the B-Experience is in this way explanatorily prior to conscious judgement about the 'A' or the 'B'. An alternative is that conscious visual experience is the more basic form of consciousness in terms of which demonstrative thought should be explained. Visual experience of the 'A' of a kind which is shared between the A-Experience and the B-Experience is just as well placed to be the pre-condition which makes thought about the 'A' possible. Similarly, visual spatial experience in general is just as well placed as distinctively attentive spatial experience, when it comes to explaining the empirical data. The

cause of feature-binding, the absence of which makes subjects unable to report locations, may be visual experience of an object's location, rather than just the specific kind of experience which distinguishes the A-Experience from the B-Experience.

Neither the demands of an account of demonstrative reference, nor the empirical data Campbell cites, favour the view that experience of the kind which distinguishes the A-Experience from the B-Experience is explanatorily prior to conscious thought. Both the account of demonstrative reference and the empirical data are consistent with my proposal that experience of this kind consists in conscious thought.

Related comments apply to Declan Smithies' (2011) argument that conscious attention causes and justifies thought about the objects of visual experience. Smithies identifies conscious attention by reference to its distinctive phenomenology. His thesis is that attention of this kind is what provides 'rational control' of thoughts and actions which exploit information from the visual system – i.e. that attention of this kind causes and justifies those thoughts, and so is explanatorily prior to them, contrary to my analysis. Yet his arguments for this thesis concern consciousness quite generally, rather than the phenomenology distinctive of attention. They establish, at most, that conscious experience of some kind is required for the rational control of thought and action. This might just as well be visual experience of a kind which is shared between the A-Experience and the B-Experience, rather than the kind of conscious attention which distinguishes one from the other.

Smithies makes two arguments. First, he argues that his thesis accommodates the data about cases of blindsight. These are cases in which subjects perform well at a visual task, yet claim to be guessing or acting randomly when they perform it (Milner and Goodale 1995; Rosetti 1998). Visual information is exploited in the control of action, yet subjects do not feel they are acting on the basis of reasons. In Smithies' interpretation, this is because 'unconscious visual information plays only a non-rational causal role in influencing the subject's inclination to guess one way rather than another' (2011: 263). According to this interpretation, conscious experience is required for the rational control of thought and action. But nothing here suggests that the same applies, more specifically, to the phenomenology distinctive of attention. Indeed the blindsight cases are standardly interpreted as cases in which *visual experience* is missing. There's no good reason to think of them as cases in which visual experience is present, while the specific phenomenology of attention is

missing. So Smithies' interpretation of these cases does not require his thesis that conscious attention causes and justifies, or is explanatorily prior to, thought and action.

Second, Smithies argues on conceptual grounds. He claims that 'a belief or action is rational if and only if it has an introspectively accessible basis in virtue of which it has what it takes to survive critical reflection' (263). The argument is roughly as follows. Rational control of thought and action is possible only where critical reflection on thought and action is possible. Critical reflection on thought and action is possible only where the objects of thought and action are available to introspection. The objects of thought and action are available to introspection only where they are objects of conscious awareness. Therefore, what provides rational control of thought and action is a form of conscious awareness.

Nothing here requires that the phenomenology distinctive of voluntary attention causes and justifies conscious thought. For example, if Smithies' argument is successful, it does not show that the phenomenal difference between the A-Experience and the B-Experience is explanatorily prior to conscious thought. Visual experience of the 'A' of a kind that's shared between the A-Experience and the B-Experience is just as well placed to be the conscious state which causes and justifies thoughts about the 'A'. Smithies' argument does not require that you experience the 'A' in a distinctive mode or modality of conscious attention, in addition to enjoying visual experience and conscious thought of the 'A'. For all the argument shows, conscious attention to what you see may consist in changes and variations in visual experience and conscious thought. More generally, we can accept that visual experience causes and justifies thoughts about its objects, and that visual experience does this only where you attend to those objects, within the context of my broadly Jamesian analysis of conscious attention.

6.5 Jamesian Attention and Openness

According to the analysis I'm proposing, conscious attention does not pose a novel or distinctive challenge to *Phenomenal Openness*. Rather, the phenomenology of attention consists in changes and variations in what visual experience represents, together with changes and variations in conscious thought. Any adequate account of the conscious mind must accommodate not only the facts about how visual experience and

conscious thought may each vary, but also the fact that we simultaneously see and think of the same objects and properties. As well as visually experiencing an object's colour, say, you sometimes think that the object has that colour. Both elements contribute to your overall conscious experience. According to the Jamesian analysis I'm proposing, an account of the conscious mind which accommodates these facts thereby accommodates the facts about conscious attention.

Exactly where this leaves *Phenomenal Openness* depends, of course, on the nature of conscious cognition. That is a huge topic in its own right, and here I will only summarize some relevant options. I'll just sketch how the Jamesian analysis interacts with some other commitments in the area, rather than trying to settle the question about *Phenomenal Openness.*

The experience of seeing something and the experience of thinking about it are qualitatively distinct. So we can generate a familiar objection to *Phenomenal Openness*, if we make assume that visual experience and the experience of thinking sometimes consist in relations to the same items. For example, we might assume that when you both see and think about a certain property of a scene, visual experience and the experience of thinking each consist in a relation to a proposition which attributes that property to the scene (McDowell 1994). This is inconsistent with *Phenomenal Openness*, since *Phenomenal Openness* entails:

Relationism

> To have a visual experience with a certain qualitative character is just to bear *R* to a certain item.

If visual experience and the experience of thinking are qualitatively distinct, yet consist in relations to the same item, then *Relationism* does not specify a sufficient condition on visual experience with a certain qualitative character. In bearing *R* to an item, you might enjoy distinctively cognitive experience rather than visual experience. So bearing *R* to a certain item is never specific enough to form a sufficient condition on visual experience. A sufficient condition requires a more specific conscious relation, as well as more specific *relata*.

Phenomenal Openness turns, then, on the premise that visual experience and the experience of thinking never consist in relations to the same item. Although this is obvious once it's made explicit, I think it's a worth pausing over. It's easy to assume, for example, that visual

experience might consist in nothing more than a relation to a proposition, while the experience of thinking sometimes consists in nothing more than a relation to the very same proposition. But this is obviously inconsistent, given that visual experience and conscious thought differ in qualitative character. Be that as it may, there are several principled, independently motivated ways to argue for the premise that visual experience and the experience of thinking do not consist in relations to the same item.

One strategy is to argue that visual experience and conscious thought consist in relations to items in altogether different categories. For instance, we might argue that while thought is a relation to a proposition, visual experience is not. Suppose you accept *Phenomenal Openness* because you accept:

Simple Relationism

> To have a visual experience with a certain qualitative character is just to bear the experiential relation to certain particulars and properties of the visible environment.

Then you are likely to think that visual experience (or at least perceptual experience) is distinctive in this respect, to be contrasted with conscious thought. For example, Campbell argues that visual experience must have this form, because intentional capacities such as the capacity for propositional thought must be explained in terms of a more primitive form of awareness (2002: 120-126).

Equally, suppose you take perceptual experience to be intentional or representational, just as thinking is. You might still have principled reasons for thinking that visual experience and thought consist in relations to items in different ontological categories. For instance, you might take thoughts to have a propositional or conceptual structure, sensitive to what we know or to our dispositions for inference, which is lacking in visual experience. So you might take thought, but not visual experience, to consist in a relation to structured propositions (Peacocke 1992: 61ff).

Once we understand the non-perceptual aspects of conscious attention in terms of conscious cognition, this general strategy becomes available as a defence of *Phenomenal Openness*. In principle, someone might defend *Phenomenal Openness* in a parallel way, even if he took conscious attention to be a *sui generis* mode of consciousness. But that

would be *ad hoc* to say the least. What reason would he have for thinking that *sui generis* attention and visual experience are relations to items in different ontological categories? By contrast, there are independent motivations for saying that visual experience and thought relate one to items of fundamentally different kinds.

However, these motivations are controversial. Some philosophers will insist that both visual experience and thought must have the same conceptual structure, given that the former justifies the latter (McDowell 1994). Others will insist that both visual experience and thought consist in relations to intentional contents, while denying that there is any such thing as a *structured* intentional content; contents are given simply by the sets of possible worlds logically compatible with them (Stalnaker 1996).

If you want to insist on either of those things, while also defending *Phenomenal Openness*, then you'll need to pursue a different strategy. You'll need to argue that visual experience and the experience of thinking consist in relations to different specific items, albeit within the same broad category.

It's sometimes suggested that the contents of thought differ in determinacy from the contents of visual experience – that the contents of visual experience are determinate, while those of thought and belief are more abstract. Accordingly, one might argue that, although visual experience and thought are both relations to propositions, they are never relations to exactly the same propositions. As I've argued, this suggestion should not be taken too far: visual experiences, like thoughts and beliefs, often represent only determinable colours and spatial properties. That leaves open the possibility of a difference of degree – the possibility that visual contents are *more* determinate than the contents of thought. But once we see that visual experience is indeterminate – that the contrast in determinacy is at most a matter of degree – it becomes implausible that the qualitative contrast between visual experience and thought could be sustained in these terms. Thinking of a shape is not like seeing a shape in low resolution. So a more radical strategy is required, if we're to argue that visual experience and the experience of thinking consist in relations to different specific items, albeit within the same broad category.

Our problem here is that we sometimes think about the very states of affairs we see. For example, if both seeing and thinking are relations to propositions, then you sometimes represent that *p* in both visual experience and thought. Now some philosophers deny that there is anything it's like to think that *p*. No one could reasonably deny that we

have conscious thoughts, or that thinking makes a contribution to the stream of consciousness. Rather, they deny that thinking that *p* has a proprietary phenomenology – that there is a kind of experience which, necessarily, you enjoy only if and only if you think that *p.* Instead, they argue, the phenomenology of thinking that *p* can be reduced to the phenomenology of other kinds of conscious state which accompany conscious thought – for example perceptual experience, occurrent emotion, sensory imagination and effort (Tye and Wright *forthcoming*; Carruthers and Veillet *forthcoming*).

This suggests a way to defend *Phenomenal Openness*, while granting that you sometimes represent that *p* in both visual experience and thought. According to this defence of *Phenomenal Openness*, when you think that *p*, you do not thereby have a cognitive experience which should be characterised in terms of the content that *p.* For example, when your visual experience represents the cat as black, and you think that the cat is black, you do not have a cognitive experience which consists in a representation of the cat as black. You may be thinking that the cat is black, but your cognitive experience does not consist in a relation to this proposition; your experience of thinking should be characterised in other terms, for example in terms of sensory imagination involving the English word 'black', a hint of fear or affection, and so on. So when you represent that *p* in both visual experience and thought, visual experience is the only experience which consists in a relation to the intentional content that *p*. You do not have two qualitatively distinct experiences which each consist in a relation to the same item.

I think this is a very challenging line to take. If *Phenomenal Openness* is to be defended, what we use to characterise your cognitive experience must be distinct from what we use to characterise your visual experience. The same proposition must not be used twice over. In particular, it's not clear to me whether the reductive approach to cognitive experience can be made convincing for the demonstrative component of thoughts about what we see. What resources, other than a proprietary phenomenology for these thoughts, could we use to characterise the kind of experience you have, when you exploit your visual experience to think about 'that object'?

The natural thing to appeal to is conscious attention to the object. But, as we've seen, the evidence suggests that conscious attention to what you see is not exhausted by variations in visual experience. And if *Phenomenal Openness* is to be defended, we cannot appeal to a *sui generis* attentive mode of consciousness. Demonstrative cognition is a form of

action. So one proposal here is that the phenomenology of demonstrative cognition is a phenomenology of action. According to Wu (2011a), the demonstrative component of thoughts about what we see has a phenomenology which can be captured in terms of action with respect an object: the action of selecting the object as an object of demonstrative cognition. We might claim, for example, that when you exploit your visual experience to think about 'that object', your cognitive experience can be characterised in terms of a proposition about your action with respect to the object. However, if *Phenomenal Openness* is to be defended, this must not be a complex proposition which includes the very proposition in terms of which we characterise your visual experience of the object: what we express by saying 'I'm selecting that object there', for example, must be captured in terms of a proposition wholly distinct from the proposition in terms of which we characterise your visual experience of the object as being where it is. I won't pursue this problem further here. For my money, the best way to defend *Phenomenal Openness* is to argue that visual experience and conscious thought consist in relations to items in different broad categories. But as I say that's controversial in ways I cannot hope to address in depth here.

I'd like to raise one final issue about *Phenomenal Openness* before moving on. *Phenomenal Openness* and its consequence *Relationism* concern the qualitative character of specifically *visual* experience. They say that having a visual experience with a certain character is just bearing the generic experiential relation *R* to a certain item. One might hope to defend a stronger thesis, concerning the qualitative character of any experience whatsoever (Tye 1995):

Global Relationism

> To have an experience with a certain qualitative character is just to bear *R* to a certain item.

This thesis faces a challenge which the weaker *Phenomenal Openness* does not face. As I said, conscious attention seems to be a matter of degree. You can attend to something more or less. Now suppose that the non-perceptual aspects of conscious attention can be captured in terms of conscious cognition. To the extent to which these aspects of conscious attention are a matter of degree, conscious cognition must be a matter of degree. But what is it to think of something more or less? If *Global*

Relationism is to be defended, this cannot consist in variations in the mode of consciousness by which you're aware of something. It must consist, rather, in variations in what you are aware of.

This, again, seems to me a very challenging position to defend. We might appeal to variations in the determinacy of cognitive experience: for example, your cognitive experience represents more or less determinable properties of things. But on the face of it, one might think concentratedly about a determinable property. For example, you might direct all of your cognitive resources to thinking about the determinable property *red*, and think only in a relatively peripheral way about the more determinate property *scarlet*. For this reason, the phenomenology of attention seems to me to constitute a more serious challenge to *Global Relationism* than it does to *Phenomenal Openness*. However, *Phenomenal Openness* has the theoretical value I explained (§6.1), independently of *Global Relationism*.

What's more, my broadly Jamesian analysis of conscious attention has further advantages, whether or not it allows us to defend *Phenomenal Openness*. First, this analysis has the advantage of parsimony. We can explain the formal and introspective data about conscious attention without postulating an additional, *sui generis* attentive mode of consciousness. Instead, we can understand conscious attention in terms of changes and variations in what visual experience represents, together with changes and variations in conscious thought. As a corollary, conscious attention poses no new challenge to *Phenomenal Openness*, over and above the challenge posed by conscious thought. But parsimony in the forms of consciousness we postulate is plausibly an advantage independently of that.

Second, the analysis explains why Moore's observation about transparency is so attractive, whether or not *Phenomenal Openness* is true. As I said, a good theory of attention and visual experience should explain why Moore's observation has proved so compelling, even if there are aspects of the experience of attentive vision, over and above the facts about what you see. Moore said that when you reflect introspectively on the experience of a certain property, you 'look through it and see nothing but' the experienced property (1903: 446). We've seen evidence that the overall experience you enjoy, when you attend voluntarily, includes more than just the visual experience of objects and their properties. The evidence suggests that there are further aspects your experience, over and above the facts about which objects and properties you experience –

aspects of the overall experience which explain, for example, the improved performance in Yeshurun et al.'s task.

My analysis of conscious attention explains why, despite this, nothing but experienced objects and properties is available to introspective reflection on the experience of seeing. When you reflect introspectively on an experience of seeing, you attend to the objects of visual experience (Martin 1998). In introspecting the experience of something's shape, say, you attend voluntarily to the shape itself. According to the Jamesian analysis I have proposed, your overall experience here includes both visual and cognitive aspects. There is, on the one hand, your visual experience of a certain shape (at a certain level of determinacy) and on the other hand your experience of selecting this shape in thought. So there are aspects of the experience over and above visual experience of the shape. But introspective reflection on this overall experience *consists in* these further aspects of experience: to reflect in this way is to turn your thoughts to the shape which you experience, and in particular to the way the shape figures in your experience. Introspective reflection on visual experience is itself a form of conscious attention to what you see. It's the aspect of an episode of conscious attention which consists in conscious cognition, rather than in visual experience. So we should not expect introspection to take this further aspect of experience as its object, and in that way to reveal it to you.

This explanation is connected, I think, with the solution to a puzzle I raised in §1.1. The puzzle was as follows. When we reflect introspectively on the character of visual experience, we take everything we find to be an aspect of the environment, an aspect of the environment which is presented or represented in visual experience. But there are, in fact, further aspects of attentive visual experience, over and above the aspects of the environment which are presented or represented in visual experience. How, then, can it be that we do not mistake what is not a presented or represented aspect of the environment for such an aspect of the environment? How can it be that the mistake about the character of visual experience does not engender mistakes about the character of the environment?

The puzzle is especially vivid if we consider an account of the phenomenology of visual attention like Block's. According to Block (2010), the phenomenology distinctive of attention to a visible colour, say, consists partly in the colour's taking on a certain 'perceived saturation', where this is a variation in how the colour appears visually, but not a

variation in what visual experience represents. Now, suppose you take everything you find when you reflect introspectively on your experience to be an aspect of the visible environment. You will then mistake the 'perceived saturation' for such an aspect of the environment.

The solution I propose is this. The further aspects of experience which conscious attention entails, over and above the facts about what is presented or represented in visual experience, consist in conscious thought. When you reflect introspectively on your experience, these aspects of experience are identical with the episode of introspective reflection. They are not *objects* of introspection, or non-representational aspects of visual experience which we notice, and mistake for aspects of the environment. Turning your thoughts to what you see involves no such mistake, because it presents you only with the aspects of your environment which visual experience represents, not with the further, cognitive aspects of your experience. As a result, even though the overall experience of attentive seeing is not exhausted by visual experience of aspects of the environment, this is consistent with *Epistemic Openness*: wherever a property figures essentially in the character of normally-occurring visual experience, believing that your environment instantiates that property is a reliable route to knowledge.

REFERENCES

Allport, A. 1993. Attention and control: Have we been asking the wrong questions? A critical review of twenty-five years. In S. Kornblum and D. Meyer (eds.), *Attention and Performance* XIV. Cambridge MA: MIT Press.

Andriessen, J. & Bouma, H. 1976. Eccentric vision: Adverse interactions between line segments. *Vision Research* 16:71-8.

Aristotle. *De Anima.* In J. Barnes (ed.), *The Complete Works of Aristotle*. Princeton: Princeton University Press. 1991.

Aristotle. *Physics.* In J. Barnes (ed.), *The Complete Works of Aristotle.* Princeton: Princeton University Press. 1991.

Armstrong, K. & Moore, T. 2007. Rapid enhancement of visual cortical response discriminability by microstimulation of the frontal eye field. *Proceedings of the National Academy of Sciences*, 104(22): 9499–9504.

Austin, J. 1962. *How To Do Things With Words*. Oxford: Clarendon Press.

Barsalou, L. 1999. Perceptual symbol systems. *Behavioral and Brain Sciences* 22: 577-660.

Barwise, J. 1981. Scenes and other situations. *Journal of Philosophy* 78 (7): 369-397.

Beck, D. & Kastner, S. 2007. Stimulus similarity modulates competitive interactions in human visual cortex. *Journal of Vision*, 7(2),19: 1-12.

Blaser, E., Sperlng, G. and Lu, Z. 1999. Measuring the amplification of attention. *Proceedings of the National Academy of Sciences* 96 (20): 11681-11686.

Block, N. 2005. Two neural correlates of consciousness. *Trends in Cognitive Sciences* 9(2): 46-52.

Block, N. 2007. Consciousness, accessibility, and the mesh between psychology and neuroscience. *Behavioral and Brain Sciences* 30: 481–548.

Block, N. 2010. Attention and mental paint. *Philosophical Issues* 20: 23-63.

Bouma, H. 1970. Interaction effects in parafoveal letter recognition. *Nature* 226:177-8.

Bredfeldt, C. & Ringach, D. 2002. Dynamics of spatial frequency tuning in macaque V1. *Journal of Neuroscience* 22(5): 1976-1984.

Brentano, F. 1874. *Psychologie vom Empirischen Standpunkte.* Leipzig: Dunker & Humblot.

Brewer, W. 2005. Perception and content. *European Journal of Philosophy*, 14(2):165-181.

Broackes, J. 2010. What do the colour-blind see? In J. Cohen & M. Matthen (eds.), *Color Ontology and Color Science* Cambridge MA: MIT.

Broadbent, D. 1958. *Perception and Communication.* New York: Pergamon Press.

Brockmole, J., Wang, R. & Irwin, D. 2002. Temporal integration between visual images and visual percepts. *Journal of Experimental Psychology: Human Perception & Performance* 28(2): 315–34.

Burge, T. 2003. Perceptual entitlement. *Philosophy and Phenomenological Research* 57(3): 503-548.

Byrne, A. & Hilbert, D. 2003. Color realism and color science. *Behavioral and Brain Sciences* 26, 3–64.

Campbell, J. 1993. A simple view of colour. In J. Haldane & C. Wright (eds.), *Reality: Representation and Projection.* Oxford: OUP.

Campbell, J. 2002. *Reference and Consciousness.* Oxford: Clarendon Press.

Campbell, J. 2007. What's the role of spatial awareness in visual perception of objects? *Mind and Language*, 28: 548-562.

Campbell, J. 2009. Consciousness and reference. In B. McLaughlin, A. Beckermann and S. Walter (eds.), *Oxford Handbook of Philosophy of Mind.* Oxford: OUP.

Carrasco, M. & McElree, B. 2001. Covert attention accelerates the rate of visual information-processing. *Proceedings of the National Academy of Sciences* 98: 5363-5367.

Carrasco, M., Fuller, S., & Ling, S. 2008. Transient attention does increase perceived contrast of suprathreshold stimuli: A reply to Prinzmetal, Long and Leonhardt (2008). *Perception & Psychophysics 70*(7): 1151–1164.

Carrasco, M., Ling, S., & Read, S. 2004. Attention alters appearance. *Nature Neuroscience* 7: 308–313.

Carruthers, P & Veillet, B. *Forthcoming.* The case against cognitive phenomenology. In T. Bayne and M. Montague (eds.) *Cognitive Phenomenology*. Oxford: OUP.

Chalmers, D. 2004. The representational character of experience. In B. Leiter (ed.), *The Future for Philosophy*. Oxford: OUP.

Churchland, P.M. 1988. Perceptual plasticity and theoretical neutrality: A reply to Jerry Fodor. *Philosophy of Science* 55 (2): 167-187

Cowan, N. 2001. The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, 24: 87–185.

Crane, T. 2008. Causation and determinable properties: on the efficacy of shape, size and colour. In J. Hohwy & J. Kallestrup (ed.s), *Being Reduced: New essays on reduction, explanation, and causation.* Oxford: OUP.

Crane, T. 2009. Is perception a propositional attitude? Philosophical Quarterly 59 (236): 452-469.

Crick, F. and Koch, C. 1998. Consciousness and Neuroscience. *Cerebral Cortex*, 8: 97–107.

Davidson, D. 1969: The individuation of events. In N. Rescher (ed.), *Essays in Honor of Carl G. Hempel.* Dordrecht: D. Reidel Publishing Company.

Dehaene, S., & Changeux, J-P. 2004. Neural mechanisms for access to consciousness. In M. Gazzaniga (ed.), *The Cognitive Neurosciences* iii. Cambridge MA.: MIT Press.

Dehaene, S., Changeux, J.-P., Naccache, L., Sackur, J. & Sergent, C. 2006. Conscious, preconscious, and subliminal processing: a testable taxonomy. *Trends in Cognitive Sciences*, 10(5): 204–11.

Dennett, D. 1978. Intentional systems. In *Brainstorms*. Montgomery VT: Bradford Books.

DeValois, R. & DeValois, K. 1988. *Spatial Vision.* Oxford: OUP.

Dickie, I. 2010. We are acquainted with ordinary things. In R. Jeshion (ed.), *New Essays on Singular Thought*. Oxford: OUP.

Dretske, F. 1993. Conscious experience. *Mind*, 102 (406): 263-283.

Dretske, F. 1995. *Naturalizing the Mind.* Cambridge MA: MIT Press.

Dretske, F. 2004. Change blindness. *Philosophical Studies*, 120: 1-18.

Driver, J. & Baylis, G. 1989. Movement and visual attention: the spotlight metaphor breaks down. *Journal of Experimental Psychology: Human Perception and Performance* 15: 448-56.

Fodor, J. 1983. *Modularity of Mind: An Essay on Faculty Psychology*. Cambridge MA: MIT Press.

Fodor, J. 1990. *A Theory of Content and Other Essays.* Cambridge MA: MIT Press.

Frazor, R., Albrecht, D., Geisler, W., & Crane, A. 2004. Visual cortex neurons of monkeys and cats: temporal dynamics of the spatial frequency response function. *Journal of Neurophysiology* 91: 2607–2627.

Fuller, S. & Carrasco, M. 2006. Exogenous attention and color perception: Performance and appearance of saturation and hue. *Vision Research* 46: 4032-4047.

Funkhouser, E. 2006. The determinable-determinate relation. *Nous* 40(3): 548-569.

Geach, P. 1957. *Mental Acts.*London: Routledge & Kegan Paul.

Gobell, J., & Carrasco, M. 2005. Attention alters the appearance of spatial frequency and gap size. *Psychological Science* 16: 644–651.

Hardin, C. 1993. *Color for Philosophers: Unweaving the Rainbow*. Indianapolis: Hackett.

Harman, G. 1990. The intrinsic quality of experience. *Philosophical Perspectives* 4: 31-52.

Hellie, B. 2005. Noise and perceptual indiscriminability. *Mind* 114(455): 481-508.

Hochstein, S. & Ahissar, M. 2002. View from the top: Hierarchies and reverse hierarchies in the visual system. *Neuron* 36: 791-804.

Huang, L. & Pashler, H. 2007. A Boolean map theory of visual attention. *Psychological Review* 114(3): 599–631.

Hume, D. 1740/1978. *A Treatise of Human Nature,* ed. R. Selby-Bigge / P. Nidditch. Oxford: OUP.

Hurley, S. and Noë, A. 2003. Neural plasticity and consciousness. *Biology and Philosophy*, 18: 131-168.

Hyde, D. 1997. From heaps and gaps to heaps of gluts. *Mind*, 106 (424) :641-660.

Jackson, F. 1977. *Perception: A Representative Theory.* Cambridge: Cambridge University Press.

Jacob, A. *ms.* Multiplicity in experience and the nature of relations.

James, W. 1890/1950. *The Principles of Psychology*, ed. G. Miller. New York: Dover Publications.

Johnson, W. 1921. *Logic* (Vol. 1). Cambridge: Cambridge University Press.

Kalderon, M. 2007. Color primitivism. *Philosophical Review* 116 (4): 563-601.

Kanisza, G. 1985. Seeing and thinking. *Acta Psychologica* 59: 23-33.

Kaplan, D. 1989. Demonstratives. In J. Almog, J. Perry & H. Wettstein (eds.), *Themes from Kaplan.* Oxford: OUP.

Kastner, S., De Weerd, P., Desimone, R. and Ungerleider, L. 1998: Mechanisms of directed attention in the human extrastriate cortex as revealed by functional MRI. *Science* 282: 108-111.

Kind, A. 2003. What's so transparent about transparency? *Philosophical Studies* 115: 225–244.

Koch, C. & Tsuchiya, N. 2006. Attention and consciousness: Two distinct brain processes. *Trends in Cognitive Sciences*, 11(1): 16-22.

Koch, C. 2004. *The Quest for Consciousness: A Neurobiological Approach.* Englewood, Colorado: Roberts and Co.

Kouider, S., Dehaene, S., Jobert, A. & Le Bihan, D. 2007. Cerebral bases of subliminal and supraliminal priming during reading. Cerebral Cortex, 17(9): 2019–29.

Lamme, V. 2004. Separate neural definitions of visual consciousness and visual attention; a case for phenomenal awareness. *Neural Networks* 17: 861–872.

Lamme, V. 2005. Independent neural definitions of visual awareness and attention. In A. Raftopoulos (ed.), *Cognitive Penetrability of Perception.* Hauppauge, New York: Nova Science Publishers.

Landman, R., Spekreijse, H. and Lamme, V. 2003. Large capacity storage of integrated objects before change blindness. *Vision Research*, 43(2): 149–64.

Leibniz, G. 1704/1981. *New Essays on Human Understanding*, ed.s P. Remnant and J. Bennett. Cambridge: Cambridge University Press.

Leibniz, G. 1714/1989. Principles of nature and grace, in his *Philosophical Essays*, ed.s R. Ariew and D. Garber. Indianapolis: Hackett Publishing Company.

Levi, D. 2008. Crowding—An essential bottleneck for object recognition: A mini-review. *Vision Research* 48: 635–654.

Lewis, D. 1973. *Counterfactuals*. Oxford: Basil Blackwell.

Liu, T., Abrams, J., & Carrasco, M. 2009. Voluntary attention enhances contrast appearance. *Psychological Science*, *20*(3): 354–362.

Liu, T., Fuller, S & Carrasco, M. 2006. Attention alters the appearance of motion coherence. *Psychonomic Bulletin & Review* 13 (6), 1091-1096.

Locke, J. 1690/1975. *An Essay Concerning Human Understanding*. Ed. P. Nidditch. Oxford: OUP.

Loftus, G. and Irwin, D. 1998. On the relations among different measures of visible and informational persistence. *Cognitive Psychology*, 35: 135–99.

Luck, S., Chelazzi, L., Hillyard, S. and Desimone, R. 1997. Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of macaque visual cortex. *Journal of Neurophysiology*, 77: 24–42.

Mack, A. & Rock, I. 1998. Inattentional blindness: perception without attention. In R. Wright (ed.), *Visual Attention.* Oxford: OUP.

Macmillan, N., & Creelman, C. 2005. *Detection Theory: A User's Guide* (2nd ed.). Mahwah, N.J.: Lawrence Erlbaum Associates.

Marc-Wogau, K. On historical explanation. *Theoria* 28: 213-233.

Martin, M. 1997. The shallows of the mind. *Proceedings of the Aristotelian Society, Supplementary Volume* 71: 75-98.

Martin, M. 1998. Setting things before the mind. In A. O'Hear (ed.), *Contemporary Issues in Philosophy of Mind.* Cambridge: Cambridge University Press.

Martin, M. 2002. The transparency of experience. *Mind and Language*, 17: 376-425.

Martin, M. 2006. On being alienated. In T. Gendler and J. Hawthorne (ed.s), *Perceptual Experience.* Oxford: OUP.

Martin, M. 2010. What's in a look? In B. Nanay (ed.), *Perceiving the World*. Oxford: OUP.

Martin, P. Lee, B., White, A., Solomon, S. & Ruettiger, L. 2001. Chromatic sensitivity of ganglion cells in the peripheral primate retina. *Letters to Nature* 410 (19): 933-936.

McDowell, J. 1994. *Mind and World.* Cambridge MA.: Harvard University Press

McDowell, J. 1998. Singular thought and the extent of inner space. In P. Pettit & J. McDowell (eds.), *Subject, Thought and Context.*Oxford: OUP.

McGinn, C. 1983. *The Subjective View: Secondary Qualities and Indexical Thoughts*. Oxford: OUP.

McGinn, C. 1999. The appearance of colour. In his *Knowledge and Reality: Selected Essays.* Oxford: Clarendon Press.

Merleau-Ponty, M. 1945/1962. *Phenomenology of Perception*, trans. C. Smith. London: Routledge & Kegan Paul.

Millikan, R.G. 1984. *Language, Thought and Other Biological Categories.* Cambridge MA.: MIT Press

Milner, D. and Goodale, M. 1995. *The Visual Brain in Action.* Oxford: OUP.

Mole, C. 2011. *Attention is Cognitive Unison.* Oxford: OUP.

Mole, C., Smithies, D. & Wu, W. (eds.) 2011. *Attention: Philosophical and Psychological Essays.* Oxford: OUP.

Montagna, B. & Carrasco, M. 2006. Transient covert attention and the perceived rate of flicker. *Journal of Vision* 6: 955–965

Moore, G. 1903. The refutation of idealism. *Mind*, New Series, 12 (48): 433-453.

Moore, T. 2006. The neurobiology of visual attention: finding sources. *Current Opinion in Neurobiology* 16: 159–165

Moorland, J. and Cruz, A. 1959. Colour perception with the peripheral retina. *Journal of Modern Optics* 6 (2): 117-151.

Moran, J. and Desimone, R. 1985. Selective attention gates visual processing in the extrastriate cortex. *Science* 229: 782–784.

Nagel, T. 1974. What is it like to be a bat? *Philosophical Review* 83 (4): 435-50.

Nanay, B. 2010. Attention and perceptual content. *Analysis* 70, No.2: 263-270.

Nickel, B. 2007. Against Intentionalism. *Philosophical Studies* 136 (3): 279–304.

Nissen, M. 1985. Accessing features and objects: Is location special? In M. Posner and O. Marin (ed.s), *Attention and Performance* XI. Hillsdale, NJ: Erlbaum.

Noë, A. 2004. *Action in Perception.* Cambridge MA: MIT Press

Noorlander, C. Koenderink, J., den Ouden, R. and Edens, B.W. 1983. Sensitivity to spatiotemporal colour contrast in the peripheral visual field. *Vision Research* 23: 1-11.

Nozick, R. 1981. *Philosophical Explanations*. Cambridge MA: Belknap Press.

O'Regan, J.K. and Noë, A. 2001. A sensorimotor approach to vision and visual consciousness. *Behavioral and Brain Sciences*, 24: 883–975.

O'Regan, J.K., Rensink, A. and Clark, J. 1999. Change-blindness as a result of 'mudsplashes'. *Nature*, 398: 34.

Palmer, S. and Schloss, K. 2010. An ecological valence theory of human color preference. *Proceedings of the National Academy of Sciences* 107 (19): 8877–8882.

Pautz, A. 2010. A simple view of consciousness. In R. Koons & G. Bealer (eds.), *The Waning of Materialism*. Oxford: OUP.

Peacocke, C. 1986. *Thoughts: An Essay on Content.* Oxford: Blackwell.

Peacocke, C. 1989. Perceptual content. In J. Almog, J. Perry & H. Wettstein (eds.), Themes from Kaplan. Oxford: OUP.

Peacocke, C. 1992. *A Study of Concepts*, Cambridge Mass.: MIT Press.

Posner, M. 1980. Orienting of attention, *Quarterly Journal of Experimental Psychology,* 32: 3–25.

Prinz, J. 2000. A neurofunctional theory of visual consciousness. *Consciousness and Cognition* 9(2): 243–59.

Prinz, J. 2003. A neurofunctional theory of consciousness. In A. Brook and K. Akins (eds.) *Philosophy and neuroscience.* Cambridge: Cambridge University Press.

Prinzmetal, W., Amiri, H., Allen, K., & Edwards, T. 1998. Phenomenology of attention: I. Color, location, orientation, and spatial frequency. *Journal of Experimental Psychology: Human Perception and Performance*, 24 (1): 261-282.

Prinzmetal, W. & Landau, A. 2008. Dissecting spatial visual attention. In V. Coltheart (ed.), *Tutorials in Visual Cognition.* Hove: Psychology Press.

Prinzmetal, W., Long, V. & Leonhardt, J. 2008. Involuntary attention and brightness contrast. *Perception & Psychophysics* 70(7): 1139-1150.

Prior, A. 1949. Determinables, determinates and determinants (Part I). *Mind* 58(229): 1-20.

Putnam, H. 1981. *Reason, Truth and History.* Cambridge: Cambridge University Press.

Ramsey, W., Stitch, S. & Garon, J. 1990. Connectionism, elimintivism, and the future of folk psychology. In J. Tomberlin (ed.), *Philosophical Perspectives* 4: 499-533.

Reynolds, J., Chelazzi, L. & Desimone, R. 1999. Competitive mechanisms subserve attention in macaque areas V2 and V4. *Journal of Neuroscience*, 19(5):1736–1753.

Robertson, L. 2004. *Space, Objects, Minds and Brains.* New York: Psychology Press.

Rosenthal, D. 2005. *Consciousness and Mind*. Oxford: OUP.

Rosetti, Y. 1998. Implicit short-lived motor representations of space in brain damaged and healthy patients. *Consciousness and Cognition* 7 (3): 520-58.

Searle, J. 1987. Indeterminacy, empiricism, and the first person. *Journal of Philosophy*, 84(3):123-146.

Searle, J. 1989. *Intentionality: An Essay in the Philosophy of Mind.* Cambridge: Cambridge University Press.

Sergent, C. and Dehaene, S. 2004. Is consciousness a gradual phenomenon? Evidence for an all-or none bifurcation during the attentional blink. *Psychological Science*, 15: 720–28.

Shapley, R. , & Perry, V. H . 1986. Cat and monkey retinal ganglion cells and their visual functional roles. Trends in Neuroscience 9: 229-35.

Siegel, S. 2006. Which properties are represented in perception? In T.S. Gendler and J. Hawthorne (eds.) *Perceptual Experience*. Oxford: Clarendon Press.

Siewert, C.1998. *The Signifance of Consciousness.* Princeton: Princeton University Press.

Simons, D. and Chabris, C. 1999. Gorillas in our midst: Sustained inattentional blindness for dynamic events. *Perception*, 28: 1059-1074.

Smith, A.D. 1990. Of primary and secondary qualities. *Philosophical Review* 99 (2): 221-254.

Smithies, D. 2011. Attention is rational-access consciousness. In C. Mole, D. Smithies and W. Wu (eds.), *Attention: Philosophical and Psychological Essays.* Oxford: OUP.

Snowdon, P. 1992. How to interpret 'direct perception'? In T. Crane (ed.), *The Contents of Experience*. Cambridge: Cambridge University Press.

Solomon, S., Lee, B., White, A., Rüttiger, L. & Martin, P. 2005. Chromatic organization of ganglion cell receptive fields in the peripheral retina. *Journal of Neuroscience* 25 (18): 4527– 4539.

Speaks, J. 2010. Attention and intentionalism. *Philosophical Quarterly* 60, No.239: 325-342.

Sperling, G. 1960. The information available in brief visual presentations. *Psychological Monographs: General and Applied*, 74 (11, Whole No. 498):1–29.

Stalnaker, R. 1996. On a defense of the hegemony of representation. *Philosophical Issues* 7: 101-108.

Stazicker, J. 2011. Attention, visual consciousness, and indeterminacy. *Mind & Language* 26 (2): 156–184.

Strawson, G. 1994. *Mental Reality*. Cambridge MA: MIT Press.

Strawson, P. 1959. *Individuals*. London: Methuen.

Strawson, P. 1979. Perception and its objects. In *Perception and Identity: Essays present to A.J. Ayer.* London: MacMillan Press.

Stroud, B. 2009. Scepticism and the senses. *European Journal of Philosophy*, 17(4): 559–570.

Travis, C. 2004. The silence of the senses. *Mind* 113 (449): 58-94.

Treisman, A. 1988. Features and objects. *Quarterly Journal of Experimental Psychology* 40A: 201-37.

Treue, S. 2000. Neural correlates of attention in primate visual cortex. *Trends in Neuroscience* 24: 295-300.

Tye, M. 1995. *Ten Problems of Consciousness*. Cambridge MA: MIT Press.

Tye, M. 2000. *Consciousness, Color, and Content.* Cambridge MA: MIT Press.

Tye, M. 2002. Representationalism and the transparency of experience. *Nous*, *36*(1): 137–151.

Tye, M. & Wright, B. *Forthcoming*. Is there a phenomenology of thought? In T. Bayne and M. Montague (eds.), *Cognitive Phenomenology*. Oxford: OUP.

van den Berg, R., Roerdink, J. and Cornelissen, F. 2007. On the generality of crowding: visual crowding in size, saturation, and hue compared to orientation. *Journal of Vision, 7(2),14: 1-11.*

Watzl, S. 2011. Attention as structuring of the stream of consciousness. In C. Mole, D. Smithies and W. Wu (eds.), Attention: Philosophical and Psychological Essays. Oxford: OUP.

White, A. 1964. *Attention*. Oxford: Basil Blackwell

Williamson, T. 1990. *Identity and Discrimination.* Oxford: Blackwell.

Williamson, T. 2002. *Knowledge and its Limits*. Oxford: OUP.

Williamson, T. 2006. Can cognition be factorised into internal and external components? In R. Stainton (ed.), *Contemporary Debates in Cognitive Science. Oxford: Blackwell*.

Wu, W. 2011a. What is conscious attention? Philosophy and Phenomenological Research 82 (1): 93: 120.

Wu, W. 2011b. Attention as selection for action. In C. Mole, D. Smithies and W. Wu (eds.), Attention: Philosophical and Psychological Essays.Oxford: OUP.

Yablo, S. 1992. Mental causation. *Philosophical Review* 101(2). 245-280.

Yablo, S. 1995. Singling out properties. *Philosophical Perspectives* 9: 477-502.

Yeshurun, Y. and Carrasco, M. 1998. Attention improves or impairs visual performance by enhancing spatial resolution. *Nature* 396: 72-75.

Yeshurun, Y., Montagna, B. and Carrasco, M. 2008. On the flexibility of sustained attention and its effects on a texture segmentation task. *Vision Research*, 48: 80–95.